



2021 ARGONNE LEADERSHIP
COMPUTING FACILITY

Argonne 
NATIONAL LABORATORY

Science





A team led by researchers from the Kavli Institute for Theoretical Physics is using Argonne's Theta supercomputer to perform radiation hydrodynamic simulations of massive stars with rotation. This image shows the structures of the outer envelope of a rotating 80 solar mass star. The color represents the gas density while the streamlines show the flow velocity. The envelope is convective around the yellow region driven by the strong radiation field from the core and the iron opacity peak. Convection transports most of the angular momentum outward and makes the photosphere rotate rapidly. Rotation also enhances the mass loss rate from the envelope.

Image: ALCF Visualization and Data Analysis Team; Yan-Fei Jiang, Center for Computational Astrophysics, Flatiron Institute; Lars Bildsten, Kavli Institute for Theoretical Physics, University of California, Santa Barbara

2020 SCIENCE CONTENTS

03 Message from ALCF Leadership

04 Argonne Leadership Computing Facility

06 About ALCF

07 ALCF Team

08 ALCF Computing Resources

10 Next-Generation Supercomputing Resources

12 Introducing Polaris

16 Paving the Way to Exascale

26 Science

28 Accessing ALCF
Resources for Science

30 2021 ALCF Science
Highlights

31 Biological Sciences

AI-Driven Drug Discovery for
SARS-CoV-2 Proteome
Arvind Ramanathan

COVID-19 Spread and Effectiveness
of Interventions
Jonathan Ozik and Charles Macal

Finding Druggable Sites in
SARS-CoV-2 Proteins Using
Molecular Dynamics and Machine
Learning
Albert Lau

Real-Time Analysis of SARS-CoV-2
Proteins
Darren Sherrell

Simulation of Viral Infection
Propagation Through Air-Travel
Ashok Srinivasan

36 Chemistry

NWChemEx: Tackling Chemical,
Materials, and Biochemical
Challenges in the Exascale Era
Theresa Windus

Stochastic A Priori Dynamics for
Complex Reactive Chemical
Environments
Ahren Jasper

38 Computer Science

DLIO: A Data-Centric Benchmark for
Scientific Deep Learning
Applications
Huihuo Zheng and Venkatram
Vishwanath

Rendezvous Algorithms for
Large-Scale Modeling and
Simulation
Steven Plimpton and Christopher
Knight

Parallel Relational Algebra for
Logical Inferencing at Scale
Sidharth Kumar

40 Earth Science

Extreme-Scale Simulations for
Advanced Seismic Ground Motion
and Hazard Modeling
Christine Goulet

Multi-Decadal, Climate-Scale
Convection-Resolving Simulations for
North America
V. Rao Kotamarthi

42 Energy Technologies

Automatic Building Energy Modeling
Joshua New

ExaSMR: Coupled Monte Carlo
Neutronics and Fluid Flow
Simulation of Small Modular
Reactors
Steve Hamilton

44 Engineering

Crystal Plasticity from First Principles
Vasily Bulatov

45 Materials Science

Constructing and Navigating
Polymorphic Landscapes of
Molecular Crystals
Alexandre Tkatchenko

Designing Catalysts under Realistic
Operating Conditions from the
Atomic Scale
Hendrik Heinz

Dynamic Compressed Sensing for
Real-Time Tomographic
Reconstruction
Robert Hovden

Experimentally Driven Automated
Machine-Learned Interatomic
Potential for a Refractory Oxide
Ganesh Sivaraman

Predictive Modeling of Nanoporous
Materials and Multiphase Systems
J. Ilja Siepmann

Towards Automated Discovery and
Optimization of Spin Defects for
Quantum Technologies
Giulia Galli

Towards Predictive Simulations of
Functional and Quantum Materials
Paul Kent

Ultrafast Control of Functional
Materials
Priya Vashishta

53 Physics

High-Fidelity Gyrokinetic Simulation
of Tokamak and ITER Edge Physics
Choong-Seock Chang

Mira's Last Journey
Katrín Heitmann

Nucleon Axial Charge with
All-Staggered Lattice QCD
Andreas Kronfeld

Simulations of Laser Experiments to
Study MHD Turbulence and
Non-Thermal Charged Particles
Petros Tzeferacos

Towards a Definitive Model of
Core-Collapse Supernova Explosions
Adam Burrows

58 ALCF Projects

MICHAEL E. PAPKA

ALCF Director

JINI RAMPRAKASH

ALCF Deputy Director

KATHERINE RILEY

ALCF Director of Science

KALYAN KUMARAN

ALCF Director of Technology

MARK FAHEY

ALCF Director of Operations

SUSAN COGHLAN

ALCF-X Project Director

MESSAGE FROM ALCF LEADERSHIP

Late last November, ALCF opened its ThetaGPU system to the broader ALCF community and started running new simulation, data, and learning workloads on a machine that had previously focused only on COVID-19 research. Now, ThetaGPU will also be heavily used in studies ranging from climate to chemistry to cosmology.

Designed for data-centric computing, ThetaGPU contains a parallel computing accelerator with new features for HPC, data, and deep learning, and brings a powerful new resource into the fold of ALCF technologies deployed to meet the evolving needs of the center's user community.

In August, ALCF began delivery of a new HPE hybrid GPU/CPU testbed supercomputer, named Polaris, that's expected to deliver approximately 44 petaflops of peak double precision performance and nearly 1.4 exaflops of theoretical AI performance when it goes online in 2022. This powerful bridge platform will enable science teams to prepare a wide range of HPC workloads we expect to see on ALCF's future Intel-HPE exascale supercomputer, Aurora.

The next generation of computational science will advance the requirements of modeling, simulation, AI, and data tasks. ThetaGPU and Polaris, as well as a host of smaller experimental systems, are all part of ALCF's ongoing plan to meet these application demands.

In the meantime, ALCF's main computing resource, Theta, continues to support some of the biggest and most challenging scientific simulation campaigns in the world. In 2021, noteworthy research carried out on Theta included C.S. Chang's ongoing INCITE project to better understand edge plasma physics in fusion reactors, with a focus on ITER; Joshua New's ALCC effort to model individual building energy use at a national scale; and Robert Hovden's ADSP project to enable researchers using DOE light sources, such as the Advanced Photon Source, to perform real-time 3D characterization of materials while an experiment is running.

ALCF exists to aid the research community in producing quality and impactful science, which is the focus of this annual Science Report. We are proud of the scientific and programmatic achievements showcased here that are leading us to a better, safer, and more technologically advanced tomorrow. And, as always, we are grateful for the skilled and dedicated staff whose efforts are helping push the boundaries of supercomputing research every day.

ARGONNE LEADERSHIP COMPUTING FACILITY


The ALCF enables breakthroughs in science and engineering by providing supercomputing resources and expertise to the research community.





Ryan Milner, ALCF HPC systems and network administration lead, installs components as part of the Theta supercomputer upgrade.

About ALCF



The ALCF's Theta system is designed to support research at the nexus of simulation, data analytics, and artificial intelligence.

The Argonne Leadership Computing Facility, a U.S. Department of Energy (DOE) Office of Science User Facility located at Argonne National Laboratory, enables breakthroughs in science and engineering by providing supercomputing resources and expertise to the research community.

ALCF computing resources—available to researchers from academia, industry, and government agencies—support large-scale computing projects aimed at solving some of the world's most complex and challenging scientific problems. Through awards of supercomputing time and support services, the ALCF enables its users to accelerate the pace of discovery and innovation across disciplines.

Supported by the DOE's Advanced Scientific Computing Research (ASCR) program, the ALCF and its partner organization, the Oak Ridge Leadership Computing Facility, operate leadership-class supercomputers that are orders of magnitude more powerful than the systems typically used for open scientific research.

ALCF Team



The ALCF's talented and diverse staff make the facility one of the world's premier centers for scientific computing.

Science

Computational scientists with multidisciplinary domain expertise work directly with ALCF users to maximize and accelerate their research efforts. In addition, the ALCF team applies broad expertise in data science, machine learning, data visualization and analysis, and mathematics to help application teams leverage ALCF resources to pursue data-driven discoveries.

With a deep knowledge of the ALCF computing environment and experience with a wide range of numerical methods, programming models, and computational approaches, staff scientists and performance engineers help researchers optimize the performance and productivity of simulation, data, and learning applications on ALCF systems.

Operations

The ALCF's HPC systems administrators manage and support all ALCF computing systems, ensuring users have stable, secure, and highly available resources to pursue their scientific goals. This includes the ALCF's production supercomputers, supporting system environments, storage systems, and network infrastructure. The team's software developers create tools to support the ALCF computing environment, including software for user account and project management, job failure analysis, and job scheduling.

User support specialists provide technical assistance to ALCF users and manage the workflows for user accounts and projects. In the business intelligence space, staff data architects assimilate and verify ALCF data to ensure accurate reporting of facility information.

Technology

The ALCF team plays a key role in designing and validating the facility's next-generation supercomputers. By collaborating with compute vendors and the performance tools community, staff members ensure the requisite programming models, tools, debuggers, and libraries are available on ALCF platforms. The team also helps manage Argonne's Joint Laboratory for System Evaluation, which houses next-generation testbeds that enable researchers to explore and prepare for emerging computing technologies.

ALCF computer scientists, performance engineers, and software engineers develop and optimize new tools and capabilities to facilitate science on the facility's current and future computing resources. This includes the deployment of scalable machine learning frameworks, in-situ visualization and analysis capabilities, data management services, workflow packages, and container technologies. In addition, the ALCF team is actively involved in programming language standardization efforts and contributes to cross-platform libraries to further enable the portability of HPC applications.

Outreach

ALCF staff members organize and participate in training events that prepare researchers for efficient use of leadership computing systems. They also participate in a wide variety of educational activities aimed at cultivating a diverse and skilled HPC community for the future. In addition, staff outreach efforts include facilitating partnerships with industry and communicating to external audiences the impactful research enabled by ALCF resources.

Supercomputing Resources

ALCF supercomputing resources support large-scale, computationally intensive projects aimed at solving some of the world's most complex and challenging scientific problems.

FEATURE	POLARIS	THETA: KNL NODES	THETA: GPU NODES	COOLEY
Purpose	Exascale Testbed	Production Supercomputer	Production Supercomputer	Data Analysis and Visualization Cluster
Architecture	HPE Apollo 6500 Gen10+	Intel-Cray XC40	NVIDIA DGX A100	Intel Haswell
Peak Performance	44 PF (double precision)	11.7 PF	3.9 PF	293 TF
Processors per Node	3rd Gen AMD EPYC	64-core, 1.3-GHz Intel Xeon Phi 7230	2 AMD EPYC 7742	2 6-core, 2.4-GHz Intel E5-2620
GPU per Node	4 NVIDIA A100 Tensor Core	—	8 NVIDIA A100 Tensor Core	NVIDIA Tesla K80
Nodes	560	4,392	24	126
Cores	560	281,088	576	1,512
Memory	280 TB (DDR4); 87.5 TB (HBM)	843 TB (DDR4); 70 TB (HBM)	24 TB (DDR4); 7.7 TB (HBM)	47 TB (DDR4); 3 TB (GDDR5)
Interconnect	HPE Slingshot 11 with Dragonfly configuration	Aries network with Dragonfly configuration	NVIDIA QM8700 InfiniBand	FDR InfiniBand
Racks	40	24	7	6

ALCF AI-Testbed

The ALCF AI-Testbed provides an infrastructure of next-generation AI-accelerator machines that allows researchers to evaluate the usability and performance of machine learning-based applications running on the systems. AI testbeds include:

Groq

Tensor Streaming Processor	Chip-to-Chip interconnect
>26 billion transistors, 14 nm	GroqWare software stack, Onnx
SambaNova DataScale	
Reconfigurable Dataflow Unit	RDU-Connect
>40 billion transistors, 7 nm	SambaFlow software stack, PyTorch
Habana Gaudi	
Tensor processing cores	Integrated 100 GbE-based interconnect
7nm	Synapse AI Software, PyTorch, Tensorflow

Graphcore MK1

Intelligent Processing Unit (IPU)	IPU-Links interconnect
1216 IPU tiles, 14 nm	Poplar software stack, PyTorch, Tensorflow
>23 billion transistors	
Cerebras CS-1	
Wafer-Scale Engine	SwarmX fabric
>800,000 processing cores	Tensorflow, PyTorch
2.6 trillion transistors, 7 nm	

Data Storage Systems

ALCF disk storage systems provide intermediate-term storage for users to access, analyze, and share computational and experimental data. Tape storage is used to archive data from completed projects.

FEATURE	EAGLE	GRAND	THETA-FS0	THETA-FS1	SWIFT	TAPE STORAGE
File System	Lustre	Lustre	Lustre	GPFS	Lustre	—
Storage System	HPE ClusterStor E1000	HPE ClusterStor E1000	HPE Sonexion L300	IBM Elastic Storage System (ESS)	All Flash Storage Array	LTO6 and LTO8 Tape Technology
Usable Capacity	100 PB	100 PB	9 PB	7.9 PB	123 TB	300 PB
Sustained Data Transfer Rate	650 GB/s	650 GB/s	240 GB/s	400 GB/s	48 GB/s	—
Disk Drives	8,480	8,480	2,300	7,260	—	—

Networking

Networking is the fabric that ties all of the ALCF's computing systems together. InfiniBand enables communication between system I/O nodes and the ALCF's various storage systems. The Production HPC SAN is built upon NVIDIA Mellanox High Data Rate (HDR) InfiniBand hardware. Two 800-port core switches provide the backbone links between 80 edge switches, yielding 1600 total available host ports, each at 200 Gbps, in a non-blocking fat-tree topology. The full bisection bandwidth of this fabric is 320 Tbps. The HPC SAN is maintained by the NVIDIA Mellanox Unified Fabric Manager (UFM), providing Adaptive Routing to avoid congestion, as well as the NVIDIA Mellanox Self-Healing Interconnect Enhancement for Intelligent Datacenters (SHIELD) resiliency system for link fault detection and recovery.

When external communications are required, Ethernet is the interconnect of choice. Remote user access, systems maintenance and management, as well as high performance data transfers are all enabled by the Local Area Network (LAN) and Wide Area Network (WAN) Ethernet infrastructure. This connectivity is built upon a combination of Extreme Networks SLX and MLXe routers and NVIDIA Mellanox Ethernet switches.

ALCF systems connect to other research institutions over multiple 100 Gbps Ethernet circuits that link to many high performance research networks, including local and regional networks like the Metropolitan Research and Education Network (MREN), as well as national and international networks like the Energy Sciences Network (ESnet) and Internet2.

Joint Laboratory for System Evaluation

Through Argonne's Joint Laboratory for System Evaluation (JLSE), the ALCF provides access to leading-edge testbeds for exploratory research aimed at evaluating future extreme-scale computing systems, technologies, and capabilities. JLSE testbeds include:

Arcticus, DevEP, Iris: Intel discrete and integrated GPU testbeds for ECP and ESP projects to develop, optimize, and scale applications and software for Aurora

Aurora Software Development Kit: Frequently updated version of the publicly available Intel oneAPI toolkit for Aurora development

Arm Ecosystem: Apollo 80 Fujitsu A64FX Arm system, NVIDIA Ampere Arm and A100 test kits, and an HPE Comanche with Marvell ARM64 CPU platform provide an ecosystem for porting applications and measuring performance on next-generation systems

Atos Quantum Learning Machine: Platform for testing and developing quantum algorithms and applications


Intel Xeon Clusters: Cascade Lake, Skylake, and Cooper Lake Xeon clusters enable a variety of research activities, including testing AI and learning applications

NVIDIA and AMD GPUs: Clusters of NVIDIA V100, A100, and A40 GPUs, and AMD MI50 and MI100 GPUs for preparing applications for heterogeneous computing architectures

Presque: Intel DAOS nodes for testing the Aurora storage system

NEXT-GENERATION SUPERCOMPUTING RESOURCES

The ALCF enters a new era of scientific computing with the arrival of the Polaris system and continued preparations for its upcoming Aurora exascale supercomputer.





Parviz Moin of Stanford University is leading an ALCC project that is using ALCF computing resources to perform high-fidelity simulations of bubble breakup and gas dissolution in turbulent oceanic environments. This image shows the generation of drops and bubbles from an oceanic breaking wave obtained from a geometric volume-of-fluid numerical simulation. The snapshot highlights the multiscale nature of wave-breaking phenomena.
Image: WH Ronald Chan, University of Colorado Boulder; Suhas S. Jain, Ali Mani, Shahab Mirjalili, Parviz Moin, and Javier Urzay, Stanford University

Introducing Polaris

The ALCF's new supercomputing resource gives researchers a testbed to prepare AI and data-intensive workloads for exascale.

With the arrival of the Polaris supercomputer at the ALCF, researchers now have a powerful new tool to prepare for science in the exascale era.

Developed in collaboration with Hewlett Packard Enterprise (HPE), Polaris is a leading-edge testbed system that will give scientists and application developers a platform to test and optimize codes for Aurora, the ALCF's upcoming Intel-HPE exascale supercomputer. Like Aurora, Polaris is a hybrid system equipped with both graphics processing units (GPUs) and central processing units (CPUs).

Aurora and DOE's other upcoming exascale machines will be capable of performing a billion billion calculations per second. As some of the world's first exascale-class supercomputers, the DOE systems will combine unprecedented processing power with advanced capabilities for artificial intelligence (AI) and data analysis, enabling researchers to tackle important scientific challenges, such as discovering new materials for clean energy applications, increasing our understanding of the global impacts of climate change, and exploring the increasingly large datasets generated at DOE experimental facilities, at a scale not possible today.

But conducting science on machines that are orders of magnitude more powerful than today's top supercomputers requires significant preparatory work. With DOE's Exascale Computing Project (ECP) and initiatives like the ALCF's Aurora Early Science Program, researchers have been working behind the scenes for years to ensure exascale applications, software, and hardware will be ready for science as soon as the first exascale systems hit the floor. Polaris, the ALCF's largest GPU-accelerated system to date, will be a valuable resource for Argonne researchers and the entire scientific community as they continue to prepare for Aurora and DOE's other upcoming exascale machines.

The HPE Apollo Gen10+ based supercomputer is equipped with 560 AMD EPYC processors and 2,240 NVIDIA A100 Tensor Core GPUs. The testbed will deliver approximately 44 petaflops of peak double precision performance and nearly 1.4 exaops of theoretical artificial intelligence (AI) performance, which is based on mixed-precision compute capabilities. Like Aurora, Polaris uses the Slingshot interconnect technology, which is designed to support the simulation, data, and machine learning workloads that will drive science in the exascale era.



The new Polaris system arrives at Argonne National Laboratory in August 2021.



Polaris, a hybrid CPU-GPU system built by HPE, will help ready scientists for the arrival of the ALCF's Aurora exascale supercomputer.

The Polaris software environment is equipped with the HPE Cray programming environment, HPE Performance Cluster Manager (HPCM) system software, and the ability to test programming models, such as OpenMP and SYCL, that will be available on Aurora and the next-generation supercomputers at the Oak Ridge Leadership Computing Facility (OLCF) and Lawrence Berkeley National Laboratory's National Energy Research Scientific Computing Center (NERSC). Polaris users will also benefit from NVIDIA's HPC software development kit, a suite of compilers, libraries, and tools for GPU code development.

The delivery and installation of Polaris began in August 2021. Initially, the testbed system will be dedicated to research teams participating in the ECP, Aurora Early Science Program, and the ALCF Data Science Program. In 2022, Polaris will be made available to the broader HPC community for a wide range of science and engineering projects.

In addition to helping pave the way to exascale, Polaris is a key resource in expanding the ALCF's scope beyond that of a traditional high-performance computing facility. With architectural features that support AI and data-centric workloads, Polaris is particularly well suited to handle the massive amounts of data being produced by large-scale simulations, light sources, telescopes, particle accelerators, and other experimental facilities. Initial efforts will be focused on integrating HPC and AI workloads with experimental facilities located at Argonne, including the Advanced Photon Source and the Center for Nanoscale Materials.

With many similarities at the system and user level, Polaris will be a key resource for researchers preparing to use the ALCF's Aurora exascale supercomputer.

FEATURE	POLARIS	AURORA
System Software	HPCM	HPCM
Programming Models	MPI, OpenMP, DPC++, Kokkos, RAJA, HIP, CUDA, OpenACC	MPI, OpenMP, DPC++, Kokkos, RAJA, HIP
Tools	CrayPat, gdb, Cray ATP, NVIDIA Nsight, cuda-gdb	CrayPat, gdb, Cray ATP, Intel VTune
MPI	CrayMPI, MPICH	CrayMPI, MPICH, Intel MPI
Multi-GPU	1 CPU : 4 GPU	2 CPU : 6 GPU
Data and Learning	DL frameworks, Cray AI stack, Python/Numba, Spark, Containers, Rapids	DL frameworks, Cray AI stack, Python/Numba, Spark, Containers, oneDAL
Math Libraries	cu* from CUDA	oneAPI

Shared features are shown in white.



Paving the Way to Exascale

ALCF staff members and application development teams are working behind the scenes to ensure Argonne's upcoming exascale system, Aurora, will propel science on Day 1.

Standing up a first-of-its-kind exascale supercomputer is a massive undertaking that requires ingenuity, diligence, and collaboration. While the sheer computational power of future exascale systems may grab the headlines, the behind-the-scenes work to prepare for their arrival continues to be a momentous feat.

With the delivery of Aurora, an Intel-Hewlett Packard Enterprise (HPE) supercomputer, drawing closer, a dedicated team of ALCF staff members has been working to ensure hardware, software, and a diverse set of scientific computing applications are ready for the research community as soon as the system is deployed for science.

The team's work covers everything from exascale code development and hardware technology evaluations to user training and close partnerships with vendors, fellow national laboratories, and DOE's Exascale Computing Project (ECP).

These activities and collaborations are laying the groundwork for Aurora to drive a new era of scientific discoveries and technological innovations at the ALCF and beyond.




Aurora's ability to handle simulation, data analysis, and artificial intelligence workloads will give researchers an unprecedented set of tools to advance scientific discovery.

Exascale Expertise

The process of planning and preparing for a new leadership-class supercomputer takes years of collaboration and coordination. It requires partnerships with vendors and the broader HPC community to test and develop various hardware and software components, validating their performance and functionality meets the needs of the scientific computing community.

The following summaries provide a look at a few of the many ALCF staff efforts underway to ready the facility and its users for the exascale computing era.



ALCF researchers Thomas Applencourt (left) and Colleen Bertoni have been instrumental in preparing compilers for Aurora.


Compiler Readiness

ALCF computational scientists Colleen Bertoni and Thomas Applencourt are working closely with the Intel compiler team to ensure the software enables code to run efficiently on the high-performance X^e processor that will power Aurora. The collaborative work involves evaluating the compiler's functionality and performance on ECP and Aurora Early Science Program (ESP) applications to verify its robustness and performance will meet the requirements of the ALCF user community.

This is the first time that Intel has implemented discrete GPU support in any of its C, C++, or Fortran compilers. When Bertoni and Applencourt began their work, the compilers were in something of a “pre-alpha” state—that is, very little formal testing had been performed on the software, which was still only functional in a somewhat rudimentary sense.

Bertoni and Applencourt work with the ECP and ESP teams to help test applications, find bugs, and identify reproducers for quality assurance testing. They then isolate and triage bugs and feature requests for Intel compiler team to address. The compiler itself is tested on a daily basis, with a full array of standard language benchmarks assessed. Partnerships with OLCF and NERSC meanwhile, help achieve consensus on key OpenMP directives for GPU compilers.

In total, they are tracking 20 applications and mini-apps—including QMCPACK, MILC, PHASTA, BerkeleyGW, WEST, and GAMESS—for testing, and tracking 250 reproducers based on bugs they have reported. Notably, Bertoni and Applencourt's efforts have spurred the development of code for math library function calls common to ALCF work, and led to Intel prioritizing implementation of the GPU version of those functions.



ALCF researchers JaeHyuk Kwack (left) and Servesh Muralidharan are helping to ensure key exascale tools and technologies are ready for the ALCF user community.

Profiling Tools

ALCF computational scientist JaeHyuk Kwack is working with Intel to ensure its performance profiling tools, Advisor and VTune, can help future Aurora users develop performant code for the system. Many important applications being developed for Aurora integrate Intel's optimized math libraries to maximize their performance; it is therefore crucial that Advisor and VTune are able to capture their performance characteristics seamlessly.

Kwack's work has included promoting a roofline analysis feature from Advisor. The roofline analysis captures application performance characteristics and subsequently determines achievable peak performance. With knowledge of these characteristics, application developers can identify performance bottlenecks in their applications and optimize code for Aurora testbed systems (and eventually for the finalized Aurora system itself). Kwack leads roofline analysis tutorials—regularly updated to reflect the most current technologies and trends—at a variety of conferences and training events.

Kwack's efforts also include work on the performance projection feature of Advisor. Intel's Ponte Vecchio GPUs are currently unavailable, but application developers still need estimates of how application performance will fare under the exascale system's architecture. Drawing on existing testbed systems to generate assessments to target the Ponte Vecchio GPUs, Advisor provides a systematic approach to estimating application performance for Aurora. Kwack has used an array of applications to provide several cases for validation of performance projection features; on his front he continues to collaborate with the Intel team to develop improved and more reliable capabilities.

Packaging and Early Hardware

ALCF computer scientist Servesh Muralidharan is coordinating efforts to improve early exascale hardware stability, making it easier for application developers to use Aurora testbeds at Argonne's Joint Laboratory for System Evaluation (JLSE).

The testbeds under Muralidharan's purview are used to develop applications that can eventually run on the Intel X^e GPU and Sapphire Rapids CPU being targeted for Aurora. Intel provides specialized driver components and software development kits (SDK), including the compilers, to run applications on early GPU silicon. These components and SDK require customization to work in the JLSE environment and be employable by application developers participating in the ECP and Aurora ESP. This is accomplished by collaborating with multiple teams at Intel.

Muralidharan's role with respect to packaging and early hardware involves building custom driver stacks and validating hardware behavior once the JLSE team installs and configures a server, followed by the challenge of building usable software stacks on top of the hardware. After a usable testbed is in place, Muralidharan helps diagnose low-level issues that arise from daily system use. These issues range from a specific code causing a hardware fault, to unexpected performance degradation resulting from driver problems. Once the problematic hardware is identified, he works with the corresponding Intel team to triage the issue and evaluate suitable patches in the JLSE testbed hardware.


His work with multiple silicon revisions of the Intel GPU testbed hardware has imparted a deeper understanding of the system's components and their interactions, such that he is able to maintain and, when necessary, reconfigure them.

Programming for Aurora

Through the DOE's Exascale Computing Project and the ALCF's Aurora Early Science Program, research teams across the country are working to prepare applications to run efficiently on the ALCF's upcoming exascale system. With access to the early exascale hardware and Aurora SDK, researchers are porting various codes, mini-apps, frameworks, and libraries to evaluate and optimize their performance using the programming models that will be supported on Aurora.

The following summaries provide a look at a handful of the many exascale application development efforts currently underway.





XGC allows researchers to simulate the edge region of magnetically confined thermonuclear fusion plasma. *Image: Dave Pugmire, Oak Ridge National Laboratory*


XGC

XGC is a gyrokinetic particle-in-cell code used to perform large-scale simulations on DOE supercomputers, and optimized for modeling edge plasma in particular. The code is the product of a consortium of researchers from academia and DOE laboratories including Argonne National Laboratory, Princeton Plasma Physics Laboratory, and Oak Ridge National Laboratory. XGC is being developed for exascale systems through efforts supported by the Aurora ESP and ECP's Whole Device Model Application (WDMapp) project, which seeks to develop a high-fidelity model of magnetically confined fusion plasmas to support ITER.

To prepare for Aurora, the XGC team is using a performance-portable approach that employs high-level, non-machine-specific libraries and programming models—Kokkos and Cabana. As a best practice for code development, the XGC team capitalizes on these efforts by employing higher-level interfaces and libraries. In so doing, they can directly benefit from the work being performed by library and programming model developers.

Moreover, without making any changes to their code, the team will be able to take advantage of two different maturing implementations of Kokkos. One implementation uses SYCL/DPC++ for GPU acceleration, and the other implementation uses OpenMP-target for acceleration. Both are expected to be broadly performance portable across architectures.

Because the application can run on any platform that supports the underlying software, the team changed from using vendor-specific programming approaches (such as OpenACC, CUDA, and Fortran) to using Kokkos and Cabana for GPU acceleration. Once the change was affected and the relevant programming layers were integrated into



NWChemEx is based on NWChem, an open-source, high-performance parallel computational chemistry code. *Image: NWChemEx Team*

the XGC code, the team achieved comparable or improved performance relative to that of vendor-specific implementations.


NWChemEx

NWChem is a widely used, open-source computational chemistry package. With the NWChemEx project supported by ECP and the Aurora ESP, a team led by researchers from Ames Laboratory and Iowa State University is working to redesign and reimplement the software package to provide a next-generation molecular modeling package for exascale.

This effort has provided the opportunity to restructure core functionality—including the elimination of longstanding bottlenecks associated with the generally successful NWChem code—concurrent with the production of sophisticated physics models intended to leverage the upcoming exascale systems. The development team is using multiple programming models, including CUDA, HIP, and DPC++ to maximize flexibility and target various hardware accelerators.

To help localize communication and thereby reduce related bottlenecks, NWChemEx is being geared such that CPUs handle communication protocols as well as any other non-intensive components (that is conditional-structure-based algorithms). Anything else—anything “embarrassingly parallel” or computationally expensive—is to be processed by GPUs.

In order to understand the degree to which the application is utilizing experimental hardware, the developers implement a multitiered analysis for tracking code performance. They then conduct a postmortem analysis to pinpoint the origin of errors and establish the scope of improvement that theoretically can be expected.



HACC is an extreme-scale cosmological code that allows researchers to study the structure and formation of the universe evolved over cosmic time. *Image: ALCF Visualization and Data Analysis Team and the HACC Team*


For Intel hardware, the developers are employing Intel's DPC++ Compatibility Tool to port any existing optimized CUDA code and translate it to DPC++. The most crucial aspect to using the Compatibility Tool is that it translates—on a timescale ranging from minutes to hours, depending on complexity—entire projects, not just mere source codes or specific functions. This two-step process—automated translation followed by manual finetuning—generates, from old CUDA code, performant DPC++ code that specifically targets Intel architectures.

HACC

HACC (Hardware/Hybrid Accelerated Cosmology Code) is a cosmological N-body and hydrodynamics simulation code designed to run at extreme scales on all DOE supercomputers. HACC is another code being developed for exascale via projects supported by the Aurora ESP and ECP.

HACC has a long history of running on GPU-based machines and on accelerated machines in general, thus providing solid groundwork for the current efforts to prepare for Aurora. The code is designed to emphasize performance over strict portability. While some key, computationally intensive sections of HACC must be rewritten for each new machine, its minimal library-dependence and the successful separation of communications kernels from compute-bound kernels allow it to quickly achieve optimal performance on most architectures.

As part of the Aurora preparations, the HACC team is testing several programming models for limited, computationally intensive sections of the code, including assessing multiple approaches to the development of DPC++. Focusing on the gravity-only kernels, the HACC team is studying two potential routes to a DPC++ version of the code. One is simply writing new DPC++ code by hand. The other is running



ECP's EQSIM project is focused on creating an unprecedented computational tool set and workflow for earthquake hazard and risk assessment. *Image: EQSIM Team*

the CUDA version of the code through a translation tool and evaluating the resulting DPC++. While this comparison process will not be undertaken for every HACC component, it can help inform the approach to porting the hydrodynamics kernels as well.

For the Aurora port, the already existing and highly performant OpenCL version of the code was a boon to the developers. In the initial setup, an OpenCL version of the code runs on Intel GPU hardware capable of verifying answers and evaluating performance. Because it can verify correct functionality, calculations, and expected levels of performance, this setup is used as a benchmark when comparing different programming models.

EQSIM/SW4

The SW4 (Seismic Waves, 4th order accuracy) application is a multidisciplinary simulation code for earthquake hazard and risk assessment being developed as part of the ECP's EQSIM project. Researchers from the ALCF and Lawrence Berkeley National Laboratory are leading an effort to use the C++ abstraction library RAJA, whose SYCL backend is currently being written.

Previous releases of SW4 were OpenMP implementations for multithreaded CPU execution. Recent releases utilize RAJA with implemented execution policies using OpenMP and CUDA statements for targeting CPUs and NVIDIA GPUs respectively. The RAJA SYCL and OpenMP-Target backends will be available for execution on Aurora. The existing execution policies will be implemented for these backends.

The porting effort was initiated with the SW4lite proxy application, which provided a development vehicle for driving preparation while also allowing the developers to quickly identify issues for rapid resolution. The RAJA-SYCL backend



As part of an Aurora ESP project, developers are preparing codes and software that will enable the ATLAS experiment to run its simulation and data analysis tasks on exascale systems. *Image: CERN*

execution policies have been implemented in the SW4lite proxy application for early testing and experimentation. Enabling the RAJA on Intel devices has been accomplished by utilizing oneAPI and several extensions in the DPC++ compiler. Intel's Unnamed kernel lambdas are critical for portability libraries to support general kernel execution. The Unified Shared Memory extension allows abstraction libraries to decouple loop execution from memory management. Intel's Extended Atomics and Global ID access have enabled support for the RAJA reduction object.

The developers have also made important use of many features of the SYCL programming model. Principal among these is the use SYCL `nd_ranges` to support fine-grained control over loop execution. The `nd_ranges` provide the flexibility required by a library to handle complex and simple loop executions. Through `nd_ranges` the RAJA-SYCL backend can launch simple one-dimensional SYCL kernels or complex three-dimensional kernels with explicit work group sizes.

ATLAS Experiment Codes

The ATLAS experiment at CERN's Large Hadron Collider (LHC) requires an immense amount of simulation for Standard Model and background modeling, as well as for general detector and upgrade studies. To this end, the developers are creating code that can be utilized on a multiplicity of architectures, including Aurora through an ALCF ESP project.

FastCaloSim, a code used for fast parametrized calorimeter simulation, has been written using CUDA, SYCL, and Kokkos, and has been run on the ALCF's Aurora testbed. The developers have refactored a large number of functions and files so as to minimize code duplication while maximizing the number of identical code paths between

the CUDA-Kokkos implementations. Consonant with their focus on diversity of architecture, the team has utilized a broad range of Kokkos backends for FastCaloSim. The developers concluded that FastCaloSim greatly underutilized GPU power, which suggested that a single GPU could be shared between multiple CPU processes, thereby reducing hardware expenses.

Ultimately, after considerable effort and amid a rapidly evolving compiler landscape, the team has run FastCaloSim successfully on each "flavor" of GPU attempted: Intel iGPUs and X^e-HP GPUs using DPC++, NVIDIA GPUs using SYCL with a CUDA backend, and AMD GPUs using hipSYCL (an implementation of SYCL over NVIDIA CUDA/AMD HIP).

Meanwhile, the developers are also implementing a Kokkos version of MadGraph, an event simulator for LHC experiments that performs particle-physics calculations to generate expected LHC-detector particle interactions. The team began with a CUDA implementation of the MadGraph algorithm and then ported it to Kokkos, which is preferred in this case because it is a third-party programming library written in C++ that enables developers to write their code in a single framework. Next, using OpenMP as the backend of a threaded parallel setup, MadGraph was deployed on an NVIDIA GPU-based setup. With CUDA as the backend, it was executed on the Intel GPU testbeds housed at the JLSE.


Because one of the developers' primary goals is to understand the limits of performance portability of various architectures, the performance of each different code implementation was compared to the others. The Kokkos version of MadGraph was able to achieve comparable performance to the CUDA and "vanilla" CUDA implementations, with metrics falling within 10 percent of each other.

Aurora Early Science Program Projects

The Aurora Early Science Program is designed to prepare key applications for the scale and architecture of the ALCF's upcoming exascale supercomputer, and field-test compilers and other software to pave the way for other production applications to run on the system.

The program is supporting five simulation projects, five data projects, and five learning projects. The diverse set of projects reflects the ALCF's effort to create an environment that supports emerging data science and machine learning approaches alongside traditional modeling and simulation-based research.

Simulation Projects



This image shows the baryon density (white) and the baryon temperature (color) of a cluster of galaxies. Image: JD Emberson and the HACC team, Argonne National Laboratory

Extending Moore's Law Computing with Quantum Monte Carlo

PI Anouar Benali
INST Argonne National Laboratory

High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

PI C.S. Chang
INST Princeton Plasma Physics Laboratory

NWChemEx: Tackling Chemical, Materials, and Biochemical Challenges in the Exascale Era

PI Theresa Windus
INST Iowa State University and Ames Laboratory


Extreme-Scale Cosmological Hydrodynamics

PI Katrin Heitmann
INST Argonne National Laboratory

Extreme-Scale Unstructured Adaptive CFD

PI Ken Jansen
INST University of Colorado Boulder

▲👤 Data Projects



The Exascale Computational Catalysis project will combine data science techniques and quantum chemistry simulations to explore the otherwise intractable phase space resulting from gas phase molecules on catalyst surfaces to find relevant configurations and the lowest transition states between them. *Image: Eric Hermes, Sandia National Laboratories*

Exascale Computational Catalysis

PI David Bross
INST Argonne National Laboratory

Dark Sky Mining

PI Salman Habib
INST Argonne National Laboratory

Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics

PI Ken Jansen
INST University of Colorado Boulder

Simulating and Learning in the ATLAS Detector at the Exascale

PI Walter Hopkins
INST Argonne National Laboratory

Extreme-Scale In-Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations

PI Amanda Randles
INST Duke University and Oak Ridge National Laboratory

▲👤 Learning Projects



The PPPL team's Fusion Recurrent Neural Network uses convolutional and recurrent neural network components to integrate both spatial and temporal information for predicting disruptions in tokamak plasmas. *Image: Julian Kates-Harbeck, Harvard University; Eliot Feibush, Princeton Plasma Physics Laboratory*

Machine Learning for Lattice Quantum Chromodynamics

PI William Detmold
INST Massachusetts Institute of Technology

Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

PI Nicola Ferrier
INST Argonne National Laboratory

Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials

PI Noa Marom
INST Carnegie Mellon University

Virtual Drug Response Prediction


PI Rick Stevens
INST Argonne National Laboratory


Accelerated Deep Learning Discovery in Fusion Energy Science

PI William Tang
INST Princeton Plasma Physics Laboratory

SCIENCE

The ALCF is accelerating scientific discoveries in many disciplines, ranging from chemistry and engineering to physics and materials science.





With help from ALCF computing resources, researchers are using an AI-driven workflow to investigate the mechanisms of infectivity of the SARS-CoV-2 spike protein. The latent space learned by the AI model provides a means to understand the conformational changes in the spike protein complex with the ACE2 receptor. *Image: Alex Brace and Arvind Ramanathan, Argonne National Laboratory; Anda Trifan, Argonne/University of Illinois at Urbana Champaign (UIUC); Lorenzo Casalino and Rommie Amaro, University of California San Diego; John Stone, UIUC*

Accessing ALCF Resources for Science

As a national user facility dedicated to open science, any researcher in the world with a large-scale computing problem can apply for time on ALCF computing resources.

Researchers gain access to ALCF systems for computational science and engineering projects through competitive, peer-reviewed allocations programs supported by the DOE and Argonne.

The ALCF also hosts competitive, peer-reviewed application programs designed to prepare key scientific applications and innovative computational methods for the architecture and scale of DOE supercomputers.

Application Programs

ADSP

The ALCF Data Science Program (ADSP) supports big data projects that require the scale and performance of leadership computing resources. ADSP projects focus on developing and improving data science techniques that will enable researchers to gain insights into very large datasets produced by experimental, simulation, or observational methods.

ESP

As part of the process of bringing a new supercomputer into production, the ALCF conducts its Early Science Program (ESP) to prepare applications for the architecture and scale of a new system. ESP projects represent a typical system workload at the ALCF and cover key scientific areas and numerical methods.

Allocation Programs

INCITE

The Innovative Novel Computational Impact on Theory and Experiment (INCITE) program aims to accelerate scientific discoveries and technological innovations by awarding ALCF computing time and resources to large-scale, computationally intensive projects that address grand challenges in science and engineering.

ALCC

The ASCR Leadership Computing Challenge (ALCC) program allocates ALCF computing resources to projects that advance the DOE mission; help to broaden the community of researchers capable of using leadership computing resources; and serve the national interests for scientific discovery, technological innovation, and economic competitiveness.

Director's Discretionary


Director's Discretionary projects are dedicated to leadership computing preparation, INCITE and ALCC scaling, and efforts to maximize scientific application efficiency and productivity on leadership computing platforms.

INCITE/ALCC BY DOMAIN

2021 INCITE

17.8M NODE HOURS


A	Biological Sciences	1 %
B	Chemistry	4
C	Computer Science	—
D	Earth Science	9
E	Energy Technologies	6
F	Engineering	15
G	Materials Science	35
H	Physics	30



2021 ALCC

7.3M NODE HOURS

A	Biological Sciences	11 %
B	Chemistry	15
C	Computer Science	1
D	Earth Science	19
E	Energy Technologies	13
F	Engineering	5
G	Materials Science	10
H	Physics	26



ALCC data are from calendar year 2021.

2021 Science Highlights

The ALCF user community pushes the boundaries of science across disciplines, advancing our knowledge of the universe at all scales.

ALCF supercomputers allow users to achieve scientific breakthroughs that would not otherwise be possible. From detailed atomic-level simulations to massive cosmological studies, researchers can investigate extremely complex physical systems and processes that are too small or large, costly, or dangerous to study in a laboratory.

Each year, ALCF users produce impressive results, whether they are developing and demonstrating novel computational methods or publishing papers in high-impact scientific journals.

In the following pages, we present a selection of notable results from projects supported by the ALCF's various allocation programs.

Among this year's highlights are a number of projects that used ALCF computing resources for COVID-19-related studies, including molecular modeling and epidemiology research aimed at accelerating the development of treatments and strategies to combat the pandemic.

You will also read about a research team using Theta to model individual building energy use at a national scale for the first time; a project employing predictive modeling and machine learning techniques to investigate promising materials for hydrogen storage applications; and several other research campaigns in diverse scientific areas, including cosmology, climate modeling, nuclear energy, and experimental data analysis.

AI-Driven Drug Discovery for SARS-CoV-2 Proteome

PI Arvind Ramanathan, Argonne National Laboratory
 AWARD Director's Discretionary
 HOURS Theta: 81,000 Node-Hours


This project seeks to address the fundamental biological mechanisms of the SARS-CoV-2 virus and associated COVID-19 disease, while simultaneously targeting the entire viral proteome to identify potential therapeutics.

CHALLENGE The researchers leverage leadership-class machines to design novel therapeutics against SARS-CoV-2 using AI approaches that integrate information from experimental observations, and rigorous, physics-based virtual screening and molecular simulations, to identify viable drugs that can inhibit viral proteins.

APPROACH The researchers, working closely with colleagues at Argonne's Advanced Photon Source, use what can be described as snapshots of the virus to determine its crystal structure. From this they try and identify sites of interest—potential targets for other molecules to bind to and/or attack. Simulations are used to provide details, including the effects of other molecules.

The AI approaches employed in this work integrate information from experimental observations, and rigorous, physics-based virtual screening and molecular simulations, to identify viable drugs that can inhibit viral proteins. These AI approaches, based on advances in deep learning and reinforcement learning, are capable of predicting how strongly a small molecule will bind to a protein as well as exploring the structural space of compounds that are predicted to bind to find more suitable variants.

RESULTS Computational screening of small molecules has resulted in identifying small molecules that can potentially inhibit viral function in wet-lab experiments. These experiments involve live human lung cell cultures being exposed to small molecules followed by subsequent measurements that monitor viral replication. The molecules



AI-driven MD simulations provide insights into how different ligands modulate the binding region of the viral ADP-ribose-1"-phosphatase protein. Ligands are shown in stick like representation and the protein is shown as a cartoon ensemble. Note that each ligand has an effect on distinct regions of the protein. *Image: Argonne National Laboratory*

are being further refined to optimize them for binding to specific viral target proteins. Using AI techniques, the team has screened millions of small molecules and is validating them at Argonne for activity against the virus.

IMPACT This work potentially could lead to the design of new generative models based on reinforcement learning for both small molecules and antibodies; and to the development of large-scale, AI-driven simulations of the entire viral particle and drugs bound to the various viral targets, as a better pathway to an antiviral drug. The team's work elucidating how the SARS-CoV-2 virus infiltrates the human immune system, detailed in a *International Journal of High Performance Computing Applications* paper, was awarded the first ACM Gordon Bell Special Prize for High Performance Computing-Based COVID-19 Research.

PUBLICATIONS

Casalino, L., A. Dommer, Z. Gaieb, E. P. Barros, T. Sztain, S.-H. Ahn, A. Trifan, A. Brace, A. Bogetti, H. Ma, H. Lee, M. Turilli, S. Khalid, L. Chong, C. Simmerling, D. J. Hardy, J. D. C. Maia, J. C. Phillips, T. Kurth, A. Stern, L. Huang, J. McCalpin, M. Tatineni, T. Gibbs, J. E. Stone, S. Jha, A. Ramanathan, and R. E. Amaro. "AI-Driven Multiscale Simulations Illuminate Mechanisms of SARS-CoV-2 Spike Dynamics," *International Journal of High Performance Computing Applications* (April 2021), SAGE Publishing.

COVID-19 Spread and Effectiveness of Interventions

PI Jonathan Ozik and Charles Macal,
Argonne National Laboratory

AWARD Director's Discretionary


HOURS Theta: 357,000 Node-Hours

This project, funded by the Joint DOE Laboratory Plan for Pandemic Modeling and Analysis Capability, oversees the development of epidemiological models to simulate the spread of COVID-19 throughout the population.

CHALLENGE Argonne researchers have developed CityCOVID, an agent-based model capable of tracking detailed COVID-19 transmission. Agent-based modeling is an approach for capturing the dynamics of heterogeneous, interacting, adaptive agents at an individual, granular level of detail. When applied to a city like Chicago, CityCOVID includes a synthetic population representing the 2.7 million residents of Chicago and the 1.2 million geo-located locations—including households, schools, workplaces, hospitals, nursing homes, dormitories, and jails—where they can co-locate. Throughout a simulated day, each agent moves from place-to-place, hour-by-hour, engaging in social activities and interactions with co-located agents, where COVID-19 exposure events can occur. The COVID-19 disease progression is modeled within each agent, including differing symptom severities, hospitalizations, and age-dependent probabilities of transitions between disease stages.

APPROACH The models pursue lines of inquiry familiar to anyone following the virus in news media—for example, the difference in outcomes that result from implementing various nonpharmaceutical interventions (NPIs) and how to safely ease off of the NPIs.

The project's significant computational demands result from the models' stochastic components, which encapsulate the underlying uncertainties and parameters of the simulation, and from the complexity of the population-level outcomes effected by the interactions of millions of individual software agents. The Argonne-developed technologies Repast HPC,



Weekly newly infected counts by ZIP code under differing individual behavior scenarios. Image: Jonathan Ozik, Argonne National Laboratory

ChiSIM (<https://github.com/Repast/chiSIM>), EMEWS, and Swift/T are used for model development and to run the large-scale parameter estimation and NPI scenario workflows on Theta.

RESULTS CityCOVID is being used to calibrate unobserved model parameters, such as the age-stratified, time-varying degree of individual protective behaviors across the population, and to simulate a variety of interventions and future scenarios, including the effects of vaccination campaigns and the impacts of variants of concern. A paper detailing these results—presented at SC20 and published in *The International Journal of High Performance Computing Applications*—was recognized as a finalist for the Gordon Bell Special Prize for High Performance Computing-Based COVID-19 Research.

IMPACT Model results have been provided throughout the pandemic to the Chicago and Illinois Departments of Public Health (the latter through the Illinois Governor's COVID-19 Modeling Task Force). The combined simulation and machine learning computing platform is being extended into a general platform for crisis decision making and planning.

PUBLICATIONS

Ozik, J., N. T. Collier, J. M. Wozniak, C. M. Macal, and M. Binois. "A Population Data-Driven Workflow for COVID-19 Modeling and Learning," *International Journal of High Performance Computing Applications* (September 2021), SAGE Publishing.


Finding Druggable Sites in SARS-CoV2 Proteins Using Molecular Dynamics and Machine Learning

PI Albert Lau, John Hopkins School of Medicine
 AWARD Director's Discretionary
 HOURS Theta: 250,000 Node-Hours

A molecular biophysics group at Johns Hopkins School of Medicine developed a joint computational and experimental approach using machine learning to accelerate novel drug discovery.

CHALLENGE Coronavirus disease 2019 (COVID-19) is caused by a novel coronavirus called Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2) and was declared a pandemic by the World Health Organization on March 11, 2020. To date, only a few repurposed drugs have shown limited benefits in critically ill patients. The challenge these researchers tackled was the ability to accurately and efficiently determine where the drug-binding sites are located on target proteins.

APPROACH The REST2 (Replica Exchange Solute Tempering –2nd generation), is a powerful sampling enhancement algorithm that accelerates infrequent conformational transitions of macromolecules by augmenting interaction energy fluctuations of a simulated system. Simulations with REST2 were critical in this project to search all important conformational transitions that require timescales beyond general simulation methodologies. Multiple copies (replicas) of each simulated protein system were generated by augmenting the interactions of the proteins to different effective temperatures. The team then applied TACTICS, a newly developed machine learning algorithm, to explore the druggability of various sites in the protein in conformations generated by REST2. MD trajectories were first processed by a clustering algorithm and then important conformations were analyzed with a random forest algorithm to identify relevant protein residues in conformations likely to bind drugs.



Druggable sites of SARS-Cov-2 MTase protein predicted by TACTICS workflow: open pockets (top) and ligand-bound pockets (bottom) with residues of drug binding sites and ligand emphasized in bold color. Image: Reproduced from *J. Chem. Inf. Model*, 61, 2897 (2021)

RESULTS TACTICS successfully identified druggable sites observed by previous experiments and refined the local residues and conformations likely to be important for binding at these sites. Moreover, TACTICS predicted several additional druggable sites.

IMPACT General oral medications are expected to present significant usage flexibility and reduced manufacturing/transportation/storage cost than antibodies for COVID-19 control. Development of a high fidelity, high-resolution all-atom simulation and modeling methodology that can predict all drug binding sites as well as their local conformations is a key step towards rational drug design. The TACTICS workflow developed here, which is capable of detecting “cryptic” binding sites that are difficult to detect without a binding ligand, opens the door for identifying potential druggable sites.

PUBLICATIONS

D.J. Evans, R.A. Yovanno, S. Rahman, D.W. Cao, M.Q. Beckett, M.H. Patel, A.F. Bandak, and A.Y.Lau, “Finding Druggable Sites in Proteins Using TACTICS”, *Journal of Chemical Information and Modeling* (June 2021), ACS.

Real-Time Analysis of SARS-CoV-2 Proteins


PI Darren Sherrell, Argonne National Laboratory
AWARD Director's Discretionary
HOURS Theta: 5,460 Node-Hours

Argonne researchers leveraged the ALCF's Theta supercomputer to analyze crystallographic images of a protein complex associated with the SARS-CoV-2 virus. The images come from Argonne's Advanced Photon Source (APS), following experiments designed to elucidate the protein complex's electrochemistry utilizing a technique known as serial synchrotron crystallography (SSX).

CHALLENGE Researchers at the APS performed experiments on a SARS-CoV-2 Nsp10+Nsp16 protein complex, providing the first low-dose, room-temperature insight into the complex's structure. Its results will give the community greater biological insight into the complex than is possible with traditional crystallography techniques. The experiment was designed to determine the metal activation of the complex and later lead to the first dynamic structural experiment of a protein related to SARS-CoV-2. Time-resolved structural dynamics help elucidate the electro-chemistry of this protein function and give insights into the virus.

APPROACH SSX experiments employ high-intensity x-rays to reveal the structures of large molecules using only fractional radiation doses compared with the requirements of traditional crystallographic techniques. The high speed of the technique leads to the generation of a vast array of data, the complexity and density of which necessitate sophisticated and computationally demanding analyses.

To support the rapid processing requirements, a team led by a team of Argonne researchers deployed an automated data acquisition, analysis, curation, and visualization pipeline, leveraging the Theta supercomputer for high-speed on-demand analysis. The pipeline reactively analyzes data as it is collected, moving images of the sample from the APS to the ALCF. The same automated pipeline then moved



Nsp10/16 surface with ligands. Researchers have developed a pipeline to connect ALCF supercomputers to APS experiments to enable real-time analysis of COVID-19 proteins, paving the way to elucidate important protein structural dynamics of the coronavirus. Image: Mateusz Wilamowski, University of Chicago, and George Minasov, Northwestern University

results to a repository and extracted metadata for publication in a data portal, which scientists can monitor during an experiment.

The pipeline generates a 256-image batch approximately every 35 seconds, with data transfers achieving speeds of 700 megabytes per second thanks to Globus, a University of Chicago-run data management service.

RESULTS As detailed in a paper published in *Proceedings of the National Academy of Sciences*, with their pipeline setup, the researchers investigated the Nsp10+Nsp16 protein complex during 2'-O methyltransferase activity using a fixed-target SSX method. Nineteen samples were analyzed across nearly 1,500 automated flows over the course of three ten-hour runs on the APS beam, during which over 700,000 images were processed on Theta. The resultant data were published to the data portal and used to further refine experimental work and configurations.

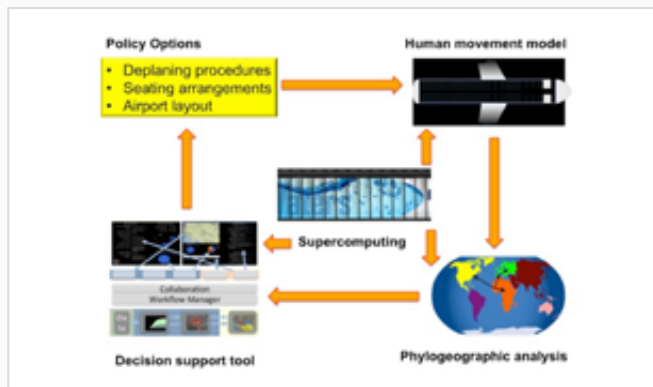
IMPACT This work paves the way to elucidate important protein structural dynamics of the coronavirus. Inhibition of Nsp16 activity may reduce viral proliferation, making this protein an attractive drug target.

PUBLICATIONS

Wilamowski, M., D. A. Sherrell, G. Minasov, Y. Kim, L. Shuvalova, A. Lavens, R. Chard, N. Maltseva, R. Jedrzejczak, M. Rosas-Lemas, N. Saint, I. T. Foster, K. Michalska, K. J. F. Satchell, and A. Joachimiak. "2'-O Methylation of RNA Cap in SARS-CoV-2 Captured by Serial Crystallography," *Proceedings of the National Academy of Sciences* (May 2021), National Academy of Sciences.

Simulation of Viral Infection Propagation Through Air-Travel

PI Ashok Srinivasan, University of West Florida
 AWARD Director's Discretionary
 HOURS Theta: 33,900 Node-Hours



Project VIPRA (Viral Infection Propagation Through Air-Travel) workflow
 Image: Ashok Srinivasan, University of West Florida

In response to the COVID-19 pandemic and the associated decline in passenger air travel, airlines introduced key changes to plane boarding and in-flight procedures to reduce infection transmission risk. However, the proximity between passengers moving in the aisle and while seated is impacted by multiple factors, such as luggage stowing, seat taking, and cabin layout. A team of researchers led by University of West Florida sought to understand the impact of new procedures on social distancing and the potential risk of disease transmission.

CHALLENGE Motivated by the idea of molecular dynamics, the researchers created two pedestrian dynamics models that simulate human walking movements by treating each pedestrian as analogous to an atom and determining contact patterns and social proximity. Because human behavior has inherent uncertainties and variations, a number of empirical parameters (such as passenger walking speed) were varied to produce numerous different scenarios. The resultant parameter space was quite large and necessitated the use of a massively parallel computer.

APPROACH The team employed a large parameter sweep to simulate 16,000 scenarios of possible passenger movement patterns, spanning several boarding processes and seating policies, to better understand the impact of policies on social distancing and risk of infection. Initial tests to evaluate convergence and determine required number of scenarios used 128 nodes of Theta for a few hours and subsequent scenario simulations required fewer resources.

RESULTS Multiple boarding procedures were examined, but only four were presented in manuscript for an Airbus 320: one zone, six zones business-first, back-to-front, and back-to-front business-first. The team examined whether the new boarding processes lead to increased or decreased social proximity, and hence risk of infection spread. Social proximity was measured via contacts between pairs of passengers whose distance was less than a specified threshold. The team also studied the reasons behind observed differences by decomposing contributions from passengers depending on whether they were seated or standing in the aisle. Simulation results showed that the back-to-front boarding policy increased social proximity by 50 percent compared to boarding procedures before the pandemic. The new policy resulted in a twofold increase compared to a one-zone policy whereby passengers are assigned seats and board randomly. Seating policies where the middle seat remains empty consistently showed a substantial reduction in social proximity during the boarding process.

IMPACT Studying the mechanisms through which social proximity is generated due to airline policies and the boarding process is crucial to reduce the risk of infection during air travel. It was found that recent changes to boarding policies by some airlines can lead to 50 percent increases in social proximity of passengers over prior boarding policies. A policy whereby passengers are assigned specific seats but board randomly is the most effective way to reduce social proximity, while maintaining empty middle seats consistently also results in substantial reductions.

PUBLICATIONS

Islam, T., M. S. Lahijani, A. Srinivasan, S. Namilae, A. Mubayi, and M. Scotch. "From Bad to Worse: Airline Boarding Changes in Response to COVID-19," *Royal Society Open Science* (April 2021), Royal Society Publishing.

Chemistry |  Simulation

NWChemEx: Tackling Chemical, Materials, and Biochemical Challenges in the Exascale Era


PI Theresa Windus, Ames Laboratory and Iowa State University
 AWARD Aurora Early Science Program and Exascale Computing Project
 HOURS Theta: 8,000 Node-Hours

This project, led by Ames Laboratory and Iowa State University, will use NWChemEx to address two challenges related to the production of advanced biofuels: the development of stress-resistant biomass feedstock and the development of catalytic processes to convert biomass-derived materials into fuels.

CHALLENGE The researchers are redesigning and reimplementing the NWChem code to enhance its scalability, performance, extensibility, and portability, positioning NWChemEx to serve as the framework for a community-wide effort to develop a comprehensive, next-generation molecular modeling package. The NWChemEx code will implement state-of-the-art algorithms for Hartree-Fock, density functional theory, and coupled cluster calculations. It will be able to effectively use multiple levels of memory as well as a partitioned global address space programming (PGAS) model to fully harness the extraordinary computational capability of the Aurora exascale computing system.

APPROACH The team's primary programming model is DPC++, a C++- and SYCL-based programming language that is part of an industry initiative to unify and simplify application development across diverse computing architectures. Intel's DPC++ Compatibility Tool helps migrate any existing CUDA code to create new DPC++ code. Further performance analysis and tuning is accomplished with Intel VTune Profiler and Intel Advisor.

To demonstrate the usefulness NWChemEx for chemical reactions, the project will probe a number of elementary chemical transformations postulated for the conversion of propanol to propene in the H-ZSM-5 zeolite. The team will run calculations using computer simulations to predict the binding energy of water and propanol and



NWChemEx will provide the understanding needed to control molecular processes underlying the production of biomass. Image: Thom H. Dunning Jr., University of Washington and Pacific Northwest National Laboratory

their reactions in the zeolite cavity to help identify appropriate biofuels.

RESULTS To help localize communication and thereby reduce related bottlenecks, NWChemEx is being designed so that CPUs handle communication protocols as well as any other non-intensive components, while GPUs will process anything computationally expensive.

The developers implement a multitiered analysis for tracking code performance so as to understand the degree to which the application is utilizing experimental hardware: roofline analysis to determine the disposition and dependencies of their algorithms, followed by actual computations performed on relevant experimental hardware to determine how efficiently processors are utilized.

IMPACT Aside from aiding the production of advanced biofuels, NWChemEx has the potential to address a number of DOE challenges, including development of next-generation batteries and new materials for solar energy conversion, simulation of chemical processes in combustion, understanding the sequestration and transport of energy byproducts in the environment, development of a science of synthesis, and design of new functional materials.

Stochastic A Priori Dynamics for Complex Reactive Chemical Environments

PI Ahren Jasper, Argonne National Laboratory

AWARD ALCC


HOURS Theta: 100,000 Node-Hours

A fundamental understanding of gas-phase chemistry is critical to the development of predictive models for a wide range of applications, including combustion, atmospheric chemistry, and low-temperature plasmas. Researchers from Argonne National Laboratory are using ALCF supercomputing resources to develop software to study the chemistry of complex reactive gas phase environments while retaining the predictive accuracy achievable in isolated single reactions.

CHALLENGE To date, high-accuracy theoretical treatments of chemical reactivity have typically involved systems with just a few to tens of atoms and rarely approach the use-inspired complexity of practical energy technologies. Computational limitations instead demand that atomistic simulations of complex environments employ approximate potential energy surfaces (PESs) and neglect nuclear quantum effects such as tunneling and electronic transitions entirely.

APPROACH With this ALCC project, researchers are applying a theoretical framework for predicting the chemistry of complex systems that is readily parallelizable and scalable. Their stochastic a priori dynamics approach is designed to enable predictive discovery in systems with use-inspired complexities by combining elementary semiclassical trajectories with ab initio two-, three-, and many-body PESs. The work makes use of the team's massively parallelized open-source codes DiNT, a feature-rich dynamics code, and PIPPy, a code for automating permutationally invariant polynomial (PIP) PES generation.

RESULTS In a study published in the *Journal of Chemical Theory and Computation*, the team detailed a strategy for constructing PIP expansions of PESs for chemical systems of any stoichiometry, demonstrating the method for pressure



Appropriate predictive gas phase theories, such as nonadiabatic semiclassical trajectories (yellow), quasiclassical trajectories (blue), master equation/transition state theory (purple), and classical dynamics (green), are used to model the collision events. Image: Ahren Jasper and Daniel Moberg, Argonne National Laboratory

dependent kinetics, three-body collisions describing transient van der Waals adducts, and nonthermal quasiclassical trajectories for calculating the reactivity of energized intermediates. Their parallelized code automates PIP generation for systems with multiple channels, enforcing flexible symmetry constraints and removing unphysical terms, all of which is required for efficiently simulating complex reactive systems that are comprised of dozens or hundreds of different elementary reaction types. This software was leveraged in two follow-up papers. In a study published in *Science*, the team determined the optimal parameters for a 16-atom PES model that shed light on the kinetics of a prototypical hydroperoxyalkyl radical ($\cdot\text{QOOH}$) intermediate. In another paper submitted for publication, the team provided databases provided databases for dictionary learning based numerical analyses designed to shrink the PIP expansion sizes and improve their efficiency.

IMPACT The team's gas-phase chemistry research will provide a foundation to develop improved predictive models for coupled reaction networks in combustion and atmospheric chemistry. The increased accuracy of a priori theoretical predictions will help in interpreting experimental results, and provide an independent source of fundamental chemical and physical information that can be used broadly throughout chemistry research.

PUBLICATIONS

Moberg, D. R., and A. W. Jasper. "Permutationally Invariant Polynomial Expansions with Unrestricted Complexity." *Journal of Chemical Theory and Computation* (September 2021), ACS.

Moberg, D. R., A. W. Jasper, and M. J. Davis. "Parsimonious potential energy surface expansions using dictionary learning with multi-pass greedy selection." *The Journal of Physical Chemistry Letters* (submitted), ACS.

Hansen, A. S., T. Bhagde, K. B. Moore, D. R. Moberg, A. W. Jasper, Y. Georgievskii, M. F. Vansco, S. J. Klippenstein, and M. I. Lester. "Watching a Hydroperoxyalkyl Radical ($\cdot\text{QOOH}$) Dissociate." *Science* (August 2021), AAAS.

DLIO: A Data-Centric Benchmark for Scientific Deep Learning Applications

PI Huihuo Zheng and Venkatram Vishwanath,
Argonne National Laboratory

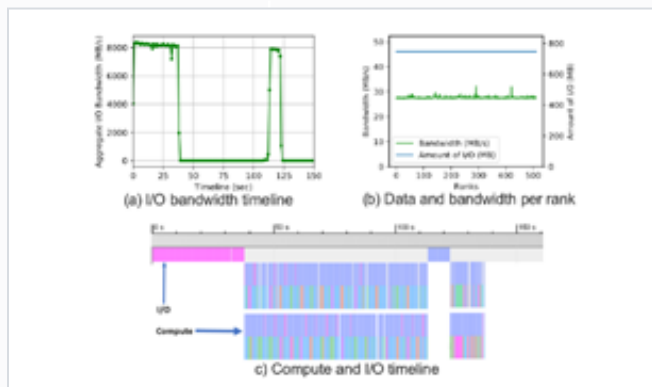
AWARD Director's Discretionary

HOURS Theta: 4,000 Node-Hours

The emergence of deep learning techniques has provided a new tool for accelerating scientific exploration and discoveries. A group of researchers from the Argonne Leadership Computing Facility (ALCF) and the Illinois Institute of Technology (IIT) set out to improve the efficiency of deep learning-driven research by developing a new benchmark, named DLIO, to investigate the data input/output (I/O) patterns of various deep learning applications.

CHALLENGE With the increase in the deep learning processing capabilities in current and future processors, the gap between computation and I/O for deep learning is expected to grow even further. The focus of current DL benchmarks has been primarily limited to understand the compute performance of deep learning applications, however, the end-to-end performance, including data processing and I/O, isn't well understood and lacking at scale. The researchers aim to provide a deeper dive into various scientific DL applications in HPC and build a representative benchmark which can further research and development.

APPROACH DLIO aims to accurately characterize the behavior of scientific DL applications and guide data-centric optimizations on modern HPC systems. To develop this, the team first characterized the behavior of modern scientific DL applications currently running on production supercomputers at Argonne Leadership Computing Facility (ALCF). In order to acquire a holistic view of how data is accessed in DL applications, the team utilized both high-level and low-level I/O profiling tools, including profilers from deep learning frameworks such as TensorFlow together with I/O profilers such as Darshan, to provide a more complete picture of the end-to-end



I/O Behavior of CANDLER NT3 DL application: Figure a) shows the aggregate bandwidth achieved of 8 GB/s for the applications. Figure b) depicts the distribution of I/O (i.e., 700 MB) and the achieved bandwidth (28 MB/s) across ranks. Figure c) is a merged timeline that shows I/O and compute do not overlap. Image: Argonne National Laboratory and Illinois Institute of Technology

application. DLIO incorporates the observed I/O behavior of these applications and is able to emulate an application's I/O performance. The benchmark suite was validated by statistically comparing the generated I/O behaviors with the application's pattern. Additionally, DLIO provides a highly tunable data-generation toolkit that can be used to project the behavior of DL applications at scale.

RESULTS The team found deep-learning applications use scientific data formats that are not well-supported by deep-learning frameworks; however, with their representative benchmark DLIO, they identified optimizations that can increase I/O efficiency by six times on existing applications.

IMPACT Using the DLIO benchmark, application developers can identify potential I/O bottlenecks in their applications on leadership systems and guide optimizations to boost the I/O performance and the overall time-to-solution. System architects can use this tool to aid the design of future systems and to evaluate the impact of various design choices on scientific machine learning applications.

PUBLICATIONS

H. Devarajan, H. Zheng, A. Kougkas, X. -H. Sun and V. Vishwanath, "DLIO: A Data-Centric Benchmark for Scientific Deep Learning Applications," 2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid), (May 2021), IEEE.

Rendezvous Algorithms for Large-Scale Modeling and Simulation


PI Christopher Knight, Argonne National Laboratory
Steven Plimpton, Sandia National Laboratories
AWARD Director's Discretionary
HOURS Mira: 625,000 Node-Hours

A variety of scientific simulation methods, when implemented in parallel on distributed-memory computers, use one or more communication patterns to send and receive data between processors. Using ALCF supercomputers, researchers explored how so-called rendezvous methods can help improve the performance of applications that use certain communication patterns on large-scale HPC systems.

CHALLENGE Rendezvous communication is a communication pattern that can be useful when processors neither know which other processors to send their data to, nor which processors will be sending them data. Often this is for the purpose of performing a parallel computation on a large number of subsets of data. Two particle simulation codes—LAMMPS and SPARTA—encountered performance bottlenecks when certain operations that have this communication pattern were run at large scale using more brute-force algorithms. To overcome this issue, researchers implemented rendezvous algorithms that provided significant performance improvements.

APPROACH The rendezvous algorithm can be implemented by defining a new partitioning of data across processors, called a rendezvous decomposition. Choosing the appropriate rendezvous decomposition enables application processors to know where to send their data, so that either a computation can be performed, or the data can be subsequently re-routed to other processors that need them. In this study, the researchers employed a general abstraction of rendezvous algorithms, where the rendezvous decomposition need not represent an alternate spatial decomposition of the data.

RESULTS As detailed in a paper published in the *Journal of Parallel and Distributed Computing*, the researchers



Intermediate rendezvous decomposition for 12 processors (red lines) overlaid on thermal and stress grids. The blue/green colored ovals are clumps of grid cells the same processor owns in the two grids. *Image: Christopher Knight, Argonne National Laboratory; Steven Plimpton, Sandia National Laboratories*

implemented rendezvous algorithms within the LAMMPS molecular dynamics code and the SPARTA direct simulation Monte Carlo code, for which some setup and other occasional operations were too slow using simpler brute-force algorithms, when running large-scale problems on supercomputers. The new rendezvous algorithms performed dramatically faster at scale. For example, LAMMPS can now enumerate bond topologies for molecular systems with billions of atoms and SPARTA can compute grid/surface intersections in models with billions of grid cells and millions of surface elements much more efficiently. The team's paper also provides an outline for a generic rendezvous pattern that is suitable for implementation in any scientific modeling code.

IMPACT This work demonstrates that rendezvous methods can help reduce performance bottlenecks for a variety of computational tasks performed in particle and grid-based codes when simpler algorithms do not scale well. It also shows that the methods can scale effectively to DOE's upcoming exascale machines.

PUBLICATIONS

Plimpton, S. J., and C. Knight. "Rendezvous Algorithms for Large-Scale Modeling and Simulation." *Journal of Parallel and Distributed Computing* (September 2020), Elsevier.

Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling

PI Christine Goulet, University of Southern California

AWARD INCITE

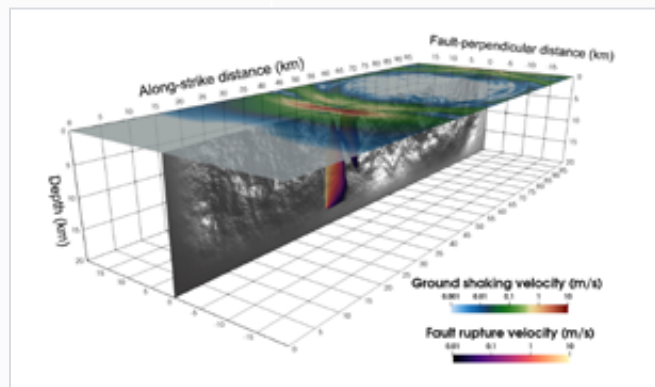
HOURS Theta: 500,000 Node-Hours

Led by the University of Southern California, this project aims to improve our understanding of earthquake system science, using state-of-the-art knowledge and software. To this end, researchers from the Southern California Earthquake Center are working to enhance their earthquake simulation and hazard mapping tools to provide the best possible information in terms of earthquake ground motion and seismic hazard.

CHALLENGE To produce realistic broadband seismograms and more accurate hazard maps, the researchers have structured their computational research plan around three objectives: (a) integrating appropriate physics models into their research software to generate realistic simulations at frequencies of engineering interest; (b) developing CyberShake workflows that integrate advanced physics capabilities for wave propagation, while ensuring its sustainability for extreme-scale computing, and (c) verifying software implementations and validating simulation products to meet criteria developed by stakeholders.

APPROACH On ALCF's Theta system, the team focuses on dynamic rupture simulations with two codes: Waveqlab3D, a FORTRAN 90 code with MPI, simulates the first-order form of the 3D elastic wave equation in collocated curvilinear grids, where the unknowns are particle velocities and stress fields; and SORD, a Fortran MPI CPU code employing the MPI-IO library that simulates spontaneous rupture within a 3D visco-plasto-elastic solid.

RESULTS The team's progress includes validation of seismograms obtained from multiphysics simulation codes—accounting for attenuation, topography, and nonlinearity—against recorded events. Theta was used to generate broadband dynamic rupture of sources incorporating geometrical fault complexity as a proxy for more



Snapshot of a magnitude 7.6 earthquake with rupture propagation and resulting wavefield at the surface computed by the dynamic rupture code SORD along a vertically dipping strike-slip fault with superimposed fault roughness. Image: Yongfei Wang, Southern California Earthquake Center, University of Southern California

detailed models of the near-field wave propagation medium. A featured research element is to calibrate the fault geometrical roughness as a dominant high-frequency radiator that can lead to comparable near-fault radiations up to 3Hz without implementing any path complexities. The resulting ground motions and fault displacements are validated against recordings and empirical models to confirm that the models capture the important physics of the problem.

IMPACT Accurate seismic hazard assessments help inform and prepare society for earthquakes, enabling the development of design and mitigation strategies that save lives and reduce economic losses. The advancement of earthquake modeling and simulation tools is critical to reducing uncertainties and improving the accuracy of seismic hazard assessments.

PUBLICATIONS

Hu, Z., D. Roten, K. B. Olsen, and S. M. Day. "Modeling of Empirical Transfer Functions with 3D Velocity Structure," *Bulletin of the Seismological Society of America* (February 2021), Seismological Society of America.

Milner, K. R., B. E. Shaw, C. A. Goulet, C. A., K. B. Richards-Dinger, S. Callaghan, T. H. Jordan, J. H., Dieterich, and E. H. Field. "A Prototype Probabilistic Seismic Hazard Model for California Constructed with Fully-Deterministic Physical Models," *Bulletin of the Seismological Society of America* (January 2021), Seismological Society of America.

O'Reilly, O., T.-Y. Yeh, K. B. Olsen, Z. Hu, A. Breuer, D. Roten, and C. A. Goulet. "A High-Order Finite-Difference Method on Staggered Curvilinear Grids for Seismic Wave Propagation Applications with Topography," *Bulletin of the Seismological Society of America* (September 2021), Seismological Society of America.

Seylabi, E., D. Restrepo, D. Asimaki, and R. Taborda. "Deterministic Ground Motion Simulations with Shallow Crust Nonlinearity at Garner Valley in Southern California," *Earthquake Engineering and Structural Dynamics* (September 2020), John Wiley and Sons.

Wang, Y., and C. A. Goulet. "Validation of Fault Displacements from Dynamic Rupture Simulations against the Observations from the 1992 Landers Earthquake," *Bulletin of the Seismological Society of America* (August 2021), Seismological Society of America.


Multi-Decadal, Climate-Scale Convection-Resolving Simulations for North America

PI V. Rao Kotamarthi, Argonne National Laboratory
 AWARD ALCC
 HOURS Theta: 300,000 Node-Hours

Regional- and local-scale climate projections are needed to help policymakers and stakeholders develop adaptation strategies and mitigation measures that address the potential effects of climate change, such as coastal flood risks, infrastructure vulnerabilities, and agricultural disruptions. To improve climate modeling capabilities, researchers from Argonne National Laboratory are adapting regional-scale models to carry out climate simulations that cover all of North America at ultra-high spatial resolution.

CHALLENGE Performing the kilometer-scale simulations needed to provide accurate climate projections at regional and local levels are beyond the limits of the current generation of climate models. To address this issue, the Argonne team is using a method known as dynamic downscaling to adapt regional-scale weather models that operate on continental scales for developing projections for North America. Their simulations focus on spatial resolutions that can resolve convection, a key parametric uncertainty in models that influences cloud and precipitation formation. A key question addressed with these model simulations is determining the causes for the observed increases in precipitation intensity increasing across the contiguous United States.

APPROACH With access to ALCF's Theta supercomputer, the Argonne team is using the state-of-the-art regional model—Weather Research and Forecasting (WRF) Model 4.3—to address this scientific challenge. Their simulations have approximately 180 million grid cells, with calculations targeted for various 20-year periods (2001-2020; 2045-2065; and 2075-2095). ALCF staff helped with the workflow by creating a dependency chain to keep simulations running as continuously as possible.



The team used convection-resolved simulations to model cloud cover at 4 km spatial resolution over North America for April 1, 2000, at a height of 4 km above surface. The detailed rendition of the cloud fields is directly comparable to data extracted from NOAA satellite datasets. *Image: V. Rao Kotamarthi, Andrea Orton, and Jiali Wang, Argonne National Laboratory*

RESULTS For the first time, the team has performed decadal-scale simulations using a regional model at ultra-high spatial resolution (4km per grid cell) that covers the North America continent. The team is currently carrying out calculations for the historical period, with other simulations planned for downscaling climate projections from coarse grid global climate model for mid and end of the century. When complete, the high-fidelity data simulated for historical and future time periods will be used to test the competing hypotheses for explaining the cause of the observed increases in precipitation intensity in the U.S. and provide an unprecedented dataset for user communities interested in climate adaptation and resiliency planning.

IMPACT The team's work is providing climate projections at spatial scales that are actionable at local and regional levels. Data from this project will be made publicly available, so decision makers can use the results to assess climate extremes and inform adaptation and resiliency strategies that combat the impacts of climate change.

Automatic Building Energy Modeling

PI Joshua New, Oak Ridge National Laboratory

AWARD ALCC


HOURS Theta: 300,000 Node-Hours

Residential and commercial buildings consume nearly three-quarters of U.S. electricity. Simulating that energy use on a broad scale can help identify ways to reduce it, cutting greenhouse gas emissions in the process. A research team from Oak Ridge National Laboratory (ORNL) is using ALCF supercomputing resources to model individual building energy use at a national scale for the first time.

CHALLENGE Creating an energy picture of a large network of buildings, while retaining actionable information at the individual building level, can illuminate areas of opportunity for deploying energy-saving technologies. However, many efforts to model buildings rely on representative prototypes of common residential and commercial buildings such as offices, warehouses, and schools. Gaps remain between what a computer model will predict and what real life will reflect in terms of energy use. To narrow those gaps, models need to be validated with actual energy use data.

APPROACH The ORNL team has developed the Automatic Building Energy Modeling (AutoBEM) software, which is used to detect buildings, generate models, and simulate building energy use for very large areas. With this ALCC project, the researchers are using the ALCF's Theta supercomputer to leverage existing organizational relationships, scalable data sources, and unique algorithms to build nation-scale building energy models.

RESULTS In a study published in *Energies*, the researchers assessed energy use across more than 178,000 buildings in Chattanooga, Tennessee. The team partnered with a municipal utility to create a digital twin of each building. The models resulting from different data sources and algorithms were compared to the utility's information on energy use for every building, down to 15-minute intervals. Then they



The Automatic Building Energy Modeling (AutoBEM) software is capable of detecting buildings, generating models, and simulating building energy use for every building in very large geographical areas. *Image: Joshua New, Oak Ridge National Laboratory*

projected eight building improvements on energy use, demand, cost, and emissions. This involved bringing all buildings up to the state's current building code and included roof insulation, lighting changes, improvements to heating and cooling efficiency, as well as more intelligent control options via smart thermostats and water heaters.

In addition, the researchers used Theta to create a publicly available data repository that includes models of 122.9 million buildings (98 percent of the nation's building stock). To do so, their AutoBEM tool leveraged DOE's open-source EnergyPlus building simulation and OpenStudio software. The models are made available by state and county so others can modify for their own purposes.

IMPACT This project is helping to identify effective energy-savings measures to create a more sustainable and resilient built environment. The team is working with companies to make the resulting building energy models and analysis free and publicly available, stimulating private sector activity towards more grid-aware energy efficiency alternatives.

PUBLICATIONS

Bass B., J. New, and W. Copeland. "Potential Energy, Demand, Emissions, and Cost Savings Distributions for Buildings in a Utility's Service Area," *Energies* (December 2020), MDPI.

New, J. R., M. Adams, B. Bass, A. Berres, and N. Clinton, Nicholas. "Model America - Data and Models of Every U.S. Building [data set]." ORNL Constellation, doi.ccs.ornl.gov/ui/doi/339 (April 2021).

Crystal Plasticity from First Principles

PI Vasily Bulatov, Lawrence Livermore National Laboratory

AWARD INCITE


HOURS Mira: 17,000,000 Node-Hours

This project, led by researchers from Lawrence Livermore National Laboratory, used large-scale molecular dynamics (MD) simulations to gain understanding of material strength and other technologically relevant mechanical properties. This research aimed to settle two longstanding controversies in classical physical metallurgy: the microscopic origin of staged strain hardening and the nature of dislocation patterns.

CHALLENGE Although it has been known since the earliest days of metallurgy that metals strengthen when mechanically deformed, no theory currently exists to predict directly from a material's intrinsic properties the mechanisms by which this hardening occurs. Large-scale MD simulations overcame the barrier to observing such mechanisms in silico, and by examining atomic-level details demonstrated that staged metal-hardening is a direct consequence of crystal rotation under uniaxial strain.

APPROACH Aluminum was chosen as a representative face-centered-cubic (fcc) metal for the current study. The LAMMPS code was used on Mira to simulate large aluminum crystals containing 300 million particles and track the system as it evolved under uniaxial strain along one of several directions. An embedded atom method model was used to describe interactions between aluminum atoms. MD trajectories ranging from 20 to 40 nanoseconds were computed, typically using the full Mira machine. Smaller simulations with 40 million particles were computed to examine six additional initial conditions.

RESULTS Dislocations were extracted from simulations every 0.1 nanoseconds for analysis and used to track population densities for each of the 12 possible slip systems



A data reduction workflow in which a detailed all-atom representation of compressed crystal (bottom) gradually morphs into a more economical representation (top) in terms of its lattice defects—dislocation lines and interfaces of twin particles. Image: Alexander Stukowski, Technische Universität Darmstadt

(combinations of symmetric layers and rows of atoms). The stress-strain response of seven aluminum single crystals subjected to uniaxial tension from simulations qualitatively matched experimental results from a classic 1975 paper. The relative ordering and shapes of the curves was reproduced despite 10 orders of magnitude difference in straining rates between simulations and experiments. Three-staged hardening was observed for all crystals that rotated under strain, while parabolic hardening without an inflection was observed for three crystals that did not rotate. All but one of tested crystals rotated as expected. The single outlying crystal that did not rotate as predicted did, however, agree with experimental observations—further evidence that the simulations and experiments were probing the same physics of crystal plasticity and that three-staged hardening is not an intrinsic material property.

IMPACT The detailed atomistic data collected from the MD simulations yielded the observation that throughout all stages of metal hardening, the basic mechanisms of this process remain constant. At odds with much of the literature, the simulations clarify that three-stage hardening is not an intrinsic material property. These simulations are important for improving the accuracy of engineering models that predict the evolution of polycrystalline textures in industrial processes like rolling, forging, and extrusion.

PUBLICATIONS

Zepeda-Ruiz, L. A., A. Stukowski, T. Oppelstrup, N. Bertin, N. R. Barton, R. Freitas, and V. V. Bulatov. "Atomistic Insights into Metal Hardening," *Nature Materials* (October 2020), Nature Publishing Group.

Constructing and Navigating Polymorphic Landscapes of Molecular Crystals

PI Alexandre Tkatchenko, University of Luxembourg


AWARD ADSP

HOURS Theta: 8,000,000 Node-Hours

Molecular crystals are ubiquitous throughout science and technology, with applications ranging from alternative energy sources to drug design. With the aim to enable the next generation of tools available for molecular materials discovery, a team led by researchers from the University of Luxembourg is exploring the polymorphic energy landscape for a substantial number of molecular crystals of interest for molecular electronics, organic photovoltaics, and pharmaceuticals.

CHALLENGE A crucial aspect of drug discovery and molecular materials design is an extensive exploration and understanding of chemical compound space (CCS)—the extremely high-dimensional space containing all feasible molecular compositions and conformations. Recently, the combination of quantum mechanical (QM) calculations with machine learning has led to considerable insight into CCS. However, progress along these lines can only happen with the availability of comprehensive QM-based datasets that adequately describe the complex structure–property relationships in molecules across CCS.

APPROACH To advance explorations into CCS, the team leveraged ALCF supercomputing resources to aid in the development of QM7-X, a comprehensive dataset of QM-based physical and chemical properties for a fundamentally important region of CCS covering small organic molecules. To do so, the researchers performed a systematic and exhaustive sampling of the (meta-)stable equilibrium structures of a subset of molecules with up to seven heavy (C, N, O, S, Cl) atoms in the GDB13 database using a density-functional tight binding approach. This was followed by the generation of 100 non-equilibrium structures displaced using normal modes for a total of approximately 4.2 million molecular structures. The team



Schematic representation of the QM7-X dataset. The “building blocks” for the QM7-X dataset are the set of ≈ 7 k molecular formulae which contain up to seven heavy (non-hydrogen) atoms from the GDB13 database. Image: Leonardo M. Sardonas, University of Luxembourg

then performed QM calculations at the PBE0+MBD level to obtain more than 42 physicochemical properties for each of molecular structures.

RESULTS In total, QM7-X contains 42 molecular (global) and atom-in-a-molecule (local) properties, which range from ground state quantities to response quantities all of which could be utilized for the construction of next-generation intra- and inter-molecular force fields. Using the tightly converged density functional theory calculations at PBE0+MBD level of theory, the properties were computed to describe the structure-property and property-property relationships in QM7-X. This level of theory has proven to be accurate and reliable for describing intramolecular degrees of freedom in addition to intermolecular interactions in organic molecular dimers, supramolecular complexes, and molecular crystals. Therefore, the calculations will be useful in validating the quality of future work using this dataset. The team detailed their findings in a paper published in *Scientific Data*.

IMPACT The QM7-X database will enable researchers to develop accurate and reliable machine learning-based techniques that will provide new insight into the complex structure–property relationships in molecules, and ultimately allow for more extensive exploration of CCS and the rational design of molecules with tailored properties.

PUBLICATIONS

Hoja, J., L. M. Sardonas, B. G. Ernst, A. Vazquez-Mayagoitia, R. A. DiStasio, and A. Tkatchenko. “QM7-X, a comprehensive dataset of quantum-mechanical properties spanning the chemical space of small organic molecules,” *Scientific Data* (February 2021), Springer Science and Business Media

Designing Catalysts under Realistic Operating Conditions from the Atomic Scale

PI Hendrik Heinz, University of Colorado Boulder


AWARD Director's Discretionary

HOURS Theta: 54,000 Node-Hours

The oxygen reduction reaction (ORR), a fundamental electrochemical reaction that reduces O_2 to water, is essential to many applications, including hydrogen fuel cells. To better understand and manage the reaction process, researchers are leveraging ALCF supercomputers to develop new computational tools and methods that are significantly more accurate than traditional methods.

CHALLENGE Platinum is a promising catalyst for ORR in fuel cells, but improvements in catalyst composition and performance are needed to make them a more cost-effective option. Due to a number of challenges in designing ORR catalysts, trial-and-error search in experiments continues to be the main mode of discovery. While many quantum mechanical and theory studies using density functional theory (DFT) have been carried out to gain insights into the ORR process, coarse approximations have presented a major roadblock to realizing quantitative predictions of reaction rates. The limitations of DFT calculations, which disregard the initial adsorption-desorption dynamics of O_2 , have led to a limited understanding of the entire ORR process.

APPROACH In a study that combines high-performance computing and experimental imaging, researchers from the University of Colorado Boulder and the University of California Los Angeles have developed a method to accurately predict the catalytic activity of platinum catalysts. The team used ALCF supercomputers to perform large-scale molecular dynamics (MD) simulations with the interface force field that reaches the necessary accuracy of 0.005 eV ($= 0.1$ kcal/mol $= 0.2$ RT) for metal-water-gas interfaces and can describe the dynamics of entire nanoscale platinum catalyst particles over tens



Engineering the atomic-scale surface features of the platinum electrode in contact with the electrolyte helps in attracting oxygen and faster conversion to water. A strongly bound oxygen molecule is highlighted in blue before the reaction. Image: Hendrik Heinz and Shiyi, University of Colorado Boulder

of nanoseconds. Using realistic initial configurations from the MD simulations, the team also performed DFT calculations to examine possible initial reaction pathways following adsorption.

RESULTS The team's simulations of electrode-electrolyte interfaces illuminated the mechanisms and predictions of the initial oxygen adsorption step of ORR. They found the relative ORR activity is determined by oxygen access to platinum surfaces, which greatly depends on specific water adlayers, while electron transfer occurs at a similar slow rate. The team's research augments our understanding of the ORR process and enables predictions of relative activities between different platinum catalysts. Their findings were published in *Science Advances*.

IMPACT The team's computational and data-intensive methods can be used to design nanostructures that maximize catalytic efficiency, as well as possible surface modifications to further optimize the cost-benefit ratio of fuel cells. The tools developed in this study can also be applied to other catalyst and electrocatalyst interfaces to enable similar advances.

PUBLICATIONS

Wang, S., E. Zhu, Y. Huang, and H. Heinz. "Direct Correlation of Oxygen Adsorption on Platinum-Electrolyte Interfaces with the Activity in the Oxygen Reduction Reaction." *Science Advances* (June 2021), AAAS

Dynamic Compressed Sensing for Real-Time Tomographic Reconstruction

PI Robert Hovden, University of Michigan

AWARD ADSP

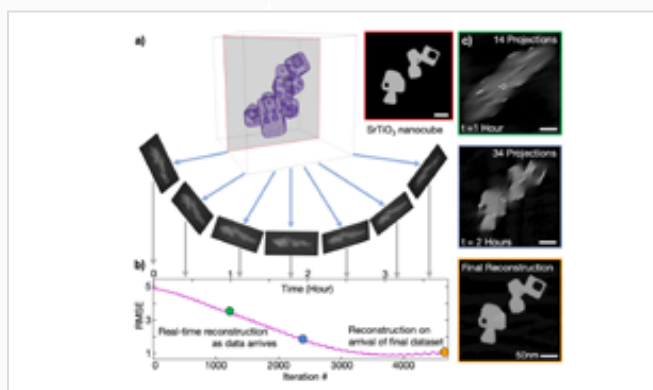
HOURS Theta: 10,000 Node-Hours

Electron and X-ray tomography allow researchers to perform 3D characterization of materials at the nano- and mesoscale, generating data that is critical to the development of nanomaterials for a wide range of applications, including solar cells and semiconductor devices. With an ALCF Data Science Program (ADSP) award, a University of Michigan-led research team is leveraging recent advancements in tomographic reconstruction algorithms, such as compressed sensing methods, to enhance and accelerate materials characterization research.

CHALLENGE Compressed sensing algorithms provide higher quality reconstructions, but they require substantially more computation time to complete, causing the rapidly expanding field of tomography to become critically bottlenecked by low throughput. To address these challenges and achieve real-time tomographic reconstruction using compressed sensing algorithms, the ADSP team has developed a dynamic framework that performs in-situ reconstruction simultaneous to data collection.

APPROACH With access to DOE supercomputing resources, the researchers are conducting comprehensive simulations for real-time electron tomography and developing reconstruction methods for through-focal tomography. The team is experimentally demonstrating their reconstruction workflow and methods on commercial scanning transmission electron microscopes and the ptychographic tomography instruments at Argonne's Advanced Photon Source (APS).

RESULTS In a paper published in *Ultramicroscopy*, the team detailed how their dynamic compressed sensing framework enables 3D specimen reconstruction in real-time as experimental data is collected. The reconstruction algorithm begins immediately upon acquiring the first projection and



a) As the tomographic experiment progresses, projections are collected across an angular range. Measured projections are fed into the dynamic CS algorithm for 3D reconstruction. b) As the amount of data increases, the root mean square error (RMSE) decreases. c) 2D slices of the 3D reconstruction at various time stamps. Image: Jonathan Schwartz, University of Michigan

dynamically updates the 3D structure as new projections arrive—unlike traditional schemes which start after the experiment is complete. This means researchers can start analysis and characterization with high-fidelity tomograms before an experiment is complete. They demonstrated that the method greatly accelerates the convergence of the total variation minimization approach and reduces the data processing waiting time by more than 70 percent.

In a study published in *Microscopy and Microanalysis*, the team investigated aberration-corrected electron microscopy, an experimental technique that can resolve the smallest atomic bond lengths in nature. The researchers carried out large-scale multislice simulations to validate the theoretical limits of the technique and demonstrate 3D atomic resolution imaging of extended specimens. With aberration-corrected electron tomography, scientists can proceed to higher resolution across larger fields-of-view to know the structure of extended specimens in all three dimensions.

IMPACT The team's approach will help advance materials characterization research by enabling real-time analysis of 3D specimens while an experiment progresses. By integrating their framework with an open-source 3D visualization and tomography software package, the team's techniques will be accessible to a wide range of researchers and enable new material characterizations across academia and industry.

PUBLICATIONS

Schwartz J., H. Zheng, M. Hanwell, Y. Jiang, and R. Hovden. "Dynamic Compressed Sensing for Real-Time Tomographic Reconstruction," *Ultramicroscopy* (December 2020). Elsevier.

Yalisove R., S. H. Sung, J. Schwartz, C. Grochner, P. Pelz, H. Zheng, Y. Jiang, C. Ophus, M. Scott, P. Ercius, and R. Hovden. "Achieving High-Resolution of Large Specimens Using Aberration-Corrected Tomography," *Microscopy and Microanalysis* (July 2020), Cambridge University Press.

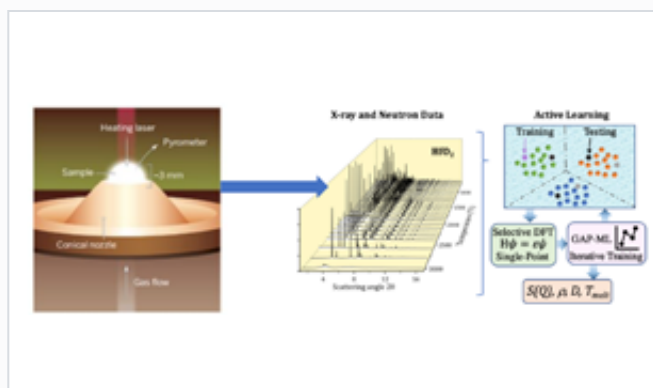
Experimentally Driven Automated Machine-Learned Interatomic Potential for a Refractory Oxide

PI Ganesh Sivaraman, Argonne National Laboratory
 AWARD Director's Discretionary
 HOURS Mira: 10,000 Node-Hours

Understanding the structure and properties of refractory oxides—ceramic materials able to withstand extreme heat—is critical for high-temperature applications. Using innovative experimental techniques and a new approach to computer simulations, a team of Argonne researchers devised a method to not only obtain precise data about the structural changes refractory oxides undergo near their melting points, but also more accurately predict other changes that can't currently be measured.

CHALLENGE In order to obtain more precise data on refractory oxides, the researchers partnered with Argonne's Advanced Photon Source (APS) and employed a beamline to examine the local and long-range structure of materials, focusing their efforts on the well-understood (and archetypally refractory) material hafnium oxide (HfO_2). The resulting data was fed into two machine-learning algorithms—one predictive and the other a filter to select the most interesting datasets—running on Argonne supercomputers. These algorithms were used to generate a new predictive model that captures qualities beyond the measurement abilities of experimentalists.

APPROACH Combining experimental and simulation techniques, the researchers' approach used an automated closed loop via an active learner initialized by x-ray and neutron diffraction measurements. This setup sequentially improved a machine-learning model until it covered an experimentally predetermined phase space. A multiphase potential was generated to provide a canonical example of HfO_2 by drawing a minimum number of training configurations from room temperature to the liquid state at approximately 2900°C. The method significantly reduced model development time and human effort. The HfO_2 itself



Experimentally synthesized structures are not necessarily the lowest-energy structures. Image: Ganesh Sivaraman, Argonne National Laboratory

was heated via aerodynamic levitation, by which procedure gasses were used to suspend small samples of the material in midair to promote isolation.

RESULTS As described in a paper published in *Physical Review Letters*, the multiphase potential created by the researchers enables prediction of a wide variety of material parameters—many of which are unavailable to experimentalists—include the ability to retain shape at high temperatures, as well as the phenomena that occur past reachable temperatures.

IMPACT Improved understanding of refractory materials helps drive innovation in materials science as well as industrial processes such as manufacturing. This work enables the study of material qualities and parameters unexplorable in the laboratory, and will be repeated on other refractory oxides.

PUBLICATIONS

Sivaraman, G., L. Gallington, A. N. Krishnamoorthy, M. Stan, G. Csányi, Á. Vázquez-Mayagoitia, and C. J. Benmore. "Experimentally Driven Automated Machine-Learned Interatomic Potential for a Refractory Oxide," *Physical Review Letters* (April 2021), APS.

Towards Automated Discovery and Optimization of Spin Defects for Quantum Technologies

PI Giulia Galli, Argonne National Laboratory and University of Chicago

AWARD INCITE


HOURS Theta: 1,200,000 Node-Hours

Spin defects in wide-bandgap semiconductors provide a promising platform for the creation of qubits, or quantum bits, the basic units of quantum information technologies. Their synthesis and optical characterization, however, present considerable challenges, and the mechanisms responsible for their generation or annihilation are poorly understood. Researchers at Argonne National Laboratory and the University of Chicago carried out atomistic and first-principles simulations to elucidate spin defect formation processes in silicon carbide (SiC) and developed a protocol, based on density functional theory calculations, to predict their photoluminescence spectra.

CHALLENGE Understanding the formation of defects in semiconductors represents a challenge for both theory and experiments. The researchers focused here on the divacancy complex (VV) in SiC, a key candidate for qubits and investigated its formation process at the atomistic level, using atomistic and quantum simulations coupled with enhanced sampling methods, and density functional theory (DFT) electronic structure calculations.

APPROACH Leveraging the Theta system to model silicon carbide at the atomistic scale, the researchers employed a combination of codes to study defects in this system. They carried out classical simulations with the LAMMPS code, and quantum simulations with the Qbox code coupled with the sampling suite of codes SSAGES, while tracking the position of vacancy defects using the VORO++ software library. They performed electronic structure calculations with both the Quantum Espresso and the Qbox codes. The research was carried out within the Midwest Integrated Center for Computational Materials.

RESULTS As detailed in a paper appearing soon in *Nature Communications*, the researchers determined that the VV



Pictorial representation of the free energy of defects in SiC, representing the formation of divacancies (left) from monovacancies (right).
Image: Elizabeth Lee, University of Chicago

formation is a thermally activated process that competes with the conversion of silicon-to-carbon monovacancies, and that VV reorientation can occur without dissociation. Moreover, their work identified pathways for the creation of spin defects consisting of antisite-double-vacancy complexes, and determined their electronic properties—potentially facilitating the realization of qubits in industrially relevant materials.

IMPACT This work lays the foundation for an integrated experimental and theoretical strategy for the design and optimization of spin defects for quantum technologies.

PUBLICATIONS

Lee, E. M. Y., A. Yu, J. J. de Pablo, and G. Galli. "Stability and Molecular Pathways to the Formation of Spin Defects in Silicon Carbide," *Nature Communications* (accepted), Springer Nature.

Predictive Modeling of Nanoporous Materials and Multiphase Systems

PI J. Ilja Siepmann, University of Minnesota

AWARD ALCC

HOURS Theta: 220,000 Node-Hours

With the ability to increase hydrogen storage capabilities at lower pressures, nanoporous materials, such as zeolites and metal-organic frameworks, can help advance the development of fuel cell vehicles. To accelerate the design and discovery of such materials, a research team led by the University of Minnesota is employing predictive hierarchical modeling and machine learning techniques on ALCF computing resources to investigate promising candidates for hydrogen storage and a variety of other energy-related applications.

CHALLENGE High-throughput molecular simulations have been the predominant approach to screening the adsorption properties of large numbers of nanoporous frameworks. Typically, a hierarchy of simulations is conducted to calculate the objective property with increasing accuracy and computational cost. These multiple rounds of simulations conclude at identifying a sufficiently small number of high-performing materials that can potentially be experimentally synthesized. To reduce the total computational cost of the high-throughput screening process and to allow for nimble variation of the objective property, machine learning models can be developed to predict adsorption properties as continuous functions of temperature, pressure, and/or composition.

APPROACH With help from the ALCF's Theta supercomputer, the University of Minnesota team developed a meta-learning method to predict gas adsorption for multiple materials over wide ranges of pressure and temperature. Meta-learning, which provides higher accuracy and improved generalization compared to fitting a model separately to each material, allows the researchers to identify the optimal hydrogen storage temperature with the highest working capacity for a given pressure difference.



A variety of CuO cluster shapes and compositions are available on the catalytic interface in reaction conditions, and the reactivity is determined by the metastable and particularly active sites on some of these cluster forms. Image: Yangzesheng Sun, University of Minnesota (icon designs from flaticon.com)

Created using data obtained from high-throughput Monte Carlo simulations of zeolites, metal-organic frameworks, and hyper-cross-linked polymers, the team's meta-learning model is able to generalize to other types of nanoporous materials and allows for few-shot learning where adsorption isotherm functions may fail because of a lack of sufficient data.

RESULTS In a study published in *Science Advances*, the team applied its meta-learning method to predict the optimal hydrogen storage temperatures for synthesized and hypothetical all-silica zeolites, hyper-cross-linked polymers, and metal-organic frameworks at four different operating conditions. The adsorption fingerprints generated by the meta-learning model show distinct features associated with nanoporous materials at the Pareto front of high-optimal temperatures and high-working capacities. The researchers also demonstrated that the predicted optimal temperature and hydrogen storage capacity given by simulation and meta-learning for a cation-exchanged zeolite are in good agreement with experimental results.

IMPACT The team's method and results provide new guidelines toward the design of hydrogen storage materials and a new route to incorporate machine learning into high-throughput materials discovery. More broadly, their work to improve the understanding and selection of nanoporous materials for energy applications could lead to significant economic and environmental benefits.

PUBLICATIONS

Sun, Y., R. F. DeJaco, Z. Li, D. Tang, S. Glante, D. S. Sholl, C. M. Colina, R. Q. Snurr, M. Thommes, M. Hartmann, and J. I. Siepmann. "Fingerprinting Diverse Nanoporous Materials for Optimal Hydrogen Storage Conditions Using Meta-Learning." *Science Advances* (July 2021), AAAS.

Towards Predictive Simulations of Functional and Quantum Materials

PI Paul Kent, Oak Ridge National Laboratory

AWARD INCITE

HOURS Theta: 1,500,000 Node-Hours


Given their unique properties and tunability, 2D semiconductor nanomaterials are good candidates for use in electronics and optical sensors. Newly developed quantum Monte Carlo (QMC) methods were used to find the structure and electronic band gap of one such material, germanium selenide (GeSe).

CHALLENGE Given the potential of 2D materials, there is great motivation to determine the properties of monolayer GeSe and how to calculate them. In bulk quantities this material forms a 3D layered structure with layers bonded by van der Waals interactions. Different density functional theory (DFT) results have yielded significantly varying geometries and band gaps.

APPROACH High-accuracy many-body diffusion Monte Carlo (DMC) methods were used to obtain an optimized GeSe structure and calculate the energy gaps of charged quasiparticle and neutral excitations. DMC was also used to verify the experimental structure and electronic properties for bulk GeSe. A newly developed DMC structural optimization method was subsequently applied to monolayer GeSe.

The researchers employed Theta for the bulk GeSe equation of state, bandgap calculations, and geometry optimization of the monolayer.

RESULTS This work resulted in the first full structural relaxation of a periodic nanostructure using QMC. Indicating the strong tunability of monolayer GeSe's optical absorption properties, the DMC energy surface has a shallow minimum at the optimal structure, while the electronic properties vary strongly with strain: not only does the magnitude of the band gap change with strain, but strain can induce a transition from a direct to an indirect gap.



Optimized geometry of GeSe monolayer using the newly developed structural optimization algorithm within Quantum Monte Carlo (colored structure) compared to the initial Density Functional Theory optimized structure (clear structure). Image: Janet Knowles, Joseph Insley, Silvio Rizzi, and Victor Mateevitsi, Argonne National Laboratory

The tunability may be a general feature of this class of materials. A multideterminant wavefunction method confirmed that potential sources of error in DMC calculations were small. The researchers also determined that no DFT exchange-correlation functional they tested could simultaneously yield both accurate band gaps and structure, indicating the importance of many-body methods such as DMC for mono- and few-layer van der Waals materials.

IMPACT GeSe is considered a promising material for light-detecting devices such as solar cells and photodetectors. Highly accurate QMC methods were used to obtain the full geometry of a complex 2D nanomaterial for the first time. The newly developed QMC-based algorithm can accurately determine the structure of a material without calculating the atomic forces. The high tunability of the band gaps indicates potential optical applications in this class of materials. This work also clearly demonstrated the need for highly accurate structural and electronic structure methods to reliably assess the properties of these materials for use in future applications.

PUBLICATIONS

Shin, H., J. T. Krogel, K. Gasperich, P. R. C. Kent, A. Benali, and O. Heinonen. "Optimized Structure and Electronic Band Gap of Monolayer GeSe from Quantum Monte Carlo Methods," *Physical Review Materials* (February 2021), APS.

Ultrafast Control of Functional Materials


PI Priya Vashishta, University of Southern California
 AWARD INCITE
 HOURS Theta: 1,600,000 Node-Hours

Artificial intelligence (AI) methods, such as machine learning and neural networks, have shown great potential in accelerating the discovery of new functional materials. With this INCITE project, researchers from the University of Southern California are combining advanced AI techniques with leadership-scale quantum dynamics simulations and experimental data to extend the frontiers of computational materials science.

CHALLENGE It can take 10-20 years or more to move a material from initial discovery to the market. This long timeline results from the empirical discovery of promising materials as well as the trial-and-error approach to identifying scalable synthesis routes for these material candidates. In recent years, an increase in available computing power and the enhanced efficiency of ab initio and AI-driven simulation software has helped accelerate large-scale computational materials screening by enabling high-throughput simulations of several tens of thousands of materials.

APPROACH The INCITE team continues to expand the physics and chemistry that can be described by their two primary codes: QXMD, a non-adiabatic quantum molecular dynamics simulation engine; and RXMD, a reactive molecular dynamics simulation engine. They are also leveraging novel methods, such as deep neural networks and reinforcement learning (RL), to advance their research into a variety of materials. In addition, the team is benefitting from synergy between their INCITE project and a related Aurora Early Science Program project, which is helping to further improve the performance of their simulation engines on Theta.

RESULTS Using ALCF computing resources, the researchers have carried out several studies to explore new



Reinforcement-learning-assisted AI agent for MoS₂ synthesis by chemical vapor deposition. Here, the agent learns to propose optimal time-dependent synthesis conditions in terms of temperature and reactant gas concentration to create defect-free MoS₂ crystal with minimal human supervision. Image: Pankaj Rajak, Argonne National Laboratory and University of Southern California

computational methods for materials discovery and development. In a paper published in *Physical Review Letters*, the team employed two neural networks to investigate the dielectric constant and its temperature dependence for liquid water. Their scalable method, which is applicable to a range of different materials and systems, computed dielectric constants that were in good agreement with experimental data.

In addition, the team used RL for two separate studies published in *npj Computational Materials*. In one paper, researchers employed an autonomous RL agent to predict the optimal synthesis conditions for the 2D quantum material, MoS₂, using chemical vapor deposition. In the other study, they successfully applied RL to generate a wide range of highly stretchable MoS₂ kirigami structures from an extremely large search space consisting of millions of structures.

IMPACT The team's use of emerging AI techniques is accelerating efforts to identify promising material compositions and phases. Ultimately, this research will help inform the discovery and synthesis of new materials engineered for targeted applications, such as batteries, catalysts, and solar cells.

PUBLICATIONS

Krishnamoorthy A., K. Nomura, N. Baradwaj, K. Shimamura, P. Rajak, A. Mishra, S. Fukushima, F. Shimojo, R. Kalia, A. Nakano, and P. Vashishta. "Dielectric Constant of Liquid Water Determined with Neural Network Quantum Molecular Dynamics," *Physical Review Letters* (May 2021), APS.

Rajak P., A. Krishnamoorthy, A. Mishra, R. Kalia, A. Nakano, and P. Vashishta. "Autonomous Reinforcement Learning Agent for Chemical Vapor Deposition Synthesis of Quantum Materials," *npj Computational Materials* (July 2021), Springer Nature.

Rajak P., B. Wang, K. Nomura, Y. Luo, A. Nakano, R. Kalia, and P. Vashishta. "Autonomous Reinforcement Learning Agent for Stretchable Kirigami Design of 2D Materials," *npj Computational Materials* (July 2021), APS.

ExaSMR: Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors

PI Steve Hamilton, Oak Ridge National Laboratory
 AWARD Aurora Early Science Program and Exascale Computing Project
 HOURS Theta: 31,000 Node-Hours

The whirls and eddies of coolant that flow around the fuel pins play a critical role in determining the reactor thermal and hydraulics performance and give much needed information to nuclear engineers about how to best design future nuclear reactor systems, both for their normal operation and for their stress tolerance.


However, understanding the physical behavior inside an operating nuclear reactor can only be accomplished via simulation due to the high-pressure, high-temperature, and radioactive environment inside a reactor core.

Large-scale high-resolution models yield more informative simulations that can help eventually build a new, intrinsically safe nuclear reactor. To this end, computational scientists and nuclear engineers at DOE laboratories are working to complete the first ever full-core pin-resolved computational fluid dynamics (CFD) model of a small modular reactor (SMR) under the DOE's Exascale Computing Project.

CHALLENGE The ultimate research objective of ExaSMR is to carry out the full-core multi-physics simulations that couple both CFD and neutron dynamics on the upcoming cutting-edge exascale supercomputers, such as the forthcoming Aurora system.

APPROACH The driver application, ENRICO, performs inline coupling of the Nek5000 CFD module with Monte Carlo (MC) through a common API that supports two MC modules—Shift, which is targeting the Frontier architecture at ORNL, and OpenMC, which is targeting the Aurora system at ANL.

RESULTS This work has resulted in the first ever full-core pin-resolved CFD simulations of SMR and application of k- τ RANS model in a GPU-oriented CFD flow-solver, NekRS. As described in a *Nuclear Engineering and Design* paper



The steady-state turbulent kinetic energy distribution in the nekRS test case. Image: Argonne National Laboratory

detailing the team's project, it was found that the mechanisms by which the coolant mixes throughout the core remain regular and relatively consistent, enabling the researchers to leverage high-fidelity simulations of the turbulent flows in a section of the core to enhance the accuracy of the core-wide computational approach.

A key aspect of the modeling of SMR fuel assemblies is the presence of spacer grids, which an important role in pressurized water reactors by creating turbulent structure and enhancing the capacity of the flow to remove heat from the fuel rods containing uranium. Instead of creating a computational grid resolving all the local geometric details, the researchers developed a mathematical mechanism to reproduce the overall impact of these structures on the coolant flow without sacrificing accuracy. The mathematical mechanism created by the researchers enables them to successfully scale up the related CFD simulations to an entire SMR core for the first time.

IMPACT This work marks an important milestone on the path to the ultimate research objective of ExaSMR and permits development of novel momentum sources to model key core structures and providing information that can help reshape how we approach the challenges in reactor designs.

PUBLICATIONS

Fang, J., D. R. Shaver, A. Tomboulides, M. Min, P. Fischer, Y.-H. Lan, R. Rahaman, P. Romano, S. Benhamadouche, Y. A. Hassan, A. Kraus, and E. Merzari. "Feasibility of Full-Core Pin Resolved CFD Simulations of Small Modular Reactor with Momentum Sources," *Nuclear Engineering and Design* (April 2021), Elsevier.

High-Fidelity Gyrokinetic Simulation of Tokamak and ITER Edge Physics


PI Choong-Seock Chang, Princeton Plasma Physics Laboratory
 AWARD INCITE
 HOURS Theta: 1,500,000 Node-Hours

This multi-year INCITE project seeks to advance the understanding of the edge plasma physics in fusion reactors, with a focus on ITER, and assemble a burning plasma experiment that can demonstrate the scientific and technological feasibility of fusion.

CHALLENGE The INCITE team is performing studies on two high-priority challenges: (1) quantifying the narrowness of the heat-flux width on the ITER divertor material plates in the high-confinement mode (H-mode) operation during tenfold energy gain operation; and (2) understanding the basic physics behind the low-to-high mode L-H transition and pedestal formation at the edge, which is necessary to achieve a tenfold energy gain in ITER.

APPROACH The researchers are using the 5D gyrokinetic particle code, XGC, on DOE leadership computing resources to address some of the most difficult plasma physics questions facing ITER. The team used this extreme-scale modeling code to solve kinetic equations for the tokamak edge by modeling plasma with a large number of particles. Predictions from XGC for the divertor heat-load width on present tokamaks agreed with experimental data within the experimental error bar, yielding narrower divertor heat-loads than what are hoped for. To find the hidden parameters, the team used a supervised machine-learning program to anchor, or direct, the extrapolation from the existing tokamak data to the future ITER data obtained from the XGC high-fidelity kinetic simulation.

RESULTS The teams' findings were published in *Physics of Plasma*. A simple extrapolation to full-power ITER would give a pessimistically narrow divertor heat-load width. However, when the same XGC code was applied to the full-power ITER, it produced over a six-times wider divertor heat-load width than the formulas developed from



A supervised machine-learning program finds the hidden kinetic parameter a/r_i and a simple predictive formula for $l_q/l_{q,Eich}$ where l_q is the divertor heat-load width normalized using the Eich regression formula. Three more ITER simulations (black symbols) have been performed to verify the validity of the new ML-found formula (black dots). Image: C.S. Chang, Princeton Plasma Physics Laboratory

a simple data extrapolation from present tokamaks. Recently, this surprising result has been confirmed by increasing the model of particles to trillions on Summit at OLCF. The new predictive formula has been successfully verified by simulating lower-power ITER plasmas on Theta.

IMPACT Establishing an accurate predictive formula for the exhaust heat-load width of future doughnut-shaped tokamak fusion reactors can allow scientists to operate ITER in a more comfortable and cost-effective way for speedier progress toward the goal of 0.5 GW of fusion power production from 50 MW of input power. A more accurate formula can also allow for more reliable design of future fusion reactors, which suffer from the limitation imposed by exhaust heat-load width on the divertor plates.

PUBLICATIONS

Chang C. S., S. Ku, R. Hager, R. M. Churchill, J. Hughes, F. Köchl, A. Loarte, V. Parail, R. A. Pitts. "Constructing a new predictive scaling formula for ITER's divertor heat-load width informed by a simulation-anchored machine learning," *Physics of Plasma* (January 2021), AIP Publishing.

Physics | 🧑🔬 Simulation, Data

Mira's Last Journey

PI Katrin Heitmann, Argonne National Laboratory

AWARD ALCC


HOURS Mira: 50,000,000 Node-Hours

This project, led by Argonne National Laboratory researchers, is one of the world's largest cosmological simulations of sufficient resolution and volume to permit the generation of detailed sky maps across multiple wavebands that are targeted towards upcoming cosmological surveys. Running on the entirety of the Mira system, the simulation modeled the lifespan of the universe to help answer some of science's deepest questions.

CHALLENGE The research team structured the simulation to begin 50 million years after the Big Bang, with conditions that agree with the most up-to-date cosmological observations. Billions of years of evolution between then and now were subsequently modeled in order to create a high-resolution model of what a large portion of the universe should look at present day.

APPROACH The project was implemented using the Hardware/Hybrid Accelerated Cosmology Code (HACC) simulation and analysis framework on the full Mira system. It evolved more than 1.24 trillion particles to resolve the cosmological structures which host faint galaxies that will be observed by the Legacy Survey of Space and Time (LSST) project when it is carried out at the Vera Rubin Observatory. Cosmological parameters chosen to be consistent with the results from the Planck satellite. Analysis outputs were generated such that synthetic galaxy catalogs could be constructed using a semi-analytic modeling approach in post-processing.

RESULTS The detailed history of the evolution of cosmological structures is now being processed to create synthetic sky maps for optical and cosmic microwave background surveys. Results have been documented in a sequence of papers and published in *The Astrophysical Journal Supplement Series*.



Dark-matter-dominated halos from a small region of the simulation. The radius of the spheres is proportional to halo mass; the dominant halo is the simulation's largest, at $\sim 6 \times 10^{15}$ solar masses. Image: ALCF Visualization and Data Analysis Team and the HACC Team

IMPACT This simulation was designed to address numerous fundamental questions in cosmology; the data produced are essential for enabling the refinement of existing predictive tools and aid the development of new models, impacting both ongoing and upcoming cosmological surveys, including the Dark Energy Spectroscopic Instrument (DESI), the LSST, SPHEREx, and the "Stage-4" ground-based cosmic microwave background experiment (CMB-S4).

PUBLICATIONS

Heitmann, K., N. Frontiere, E. Rangel, P. Larsen, A. Pope, I. Sultan, T. Uram, S. Habib, H. Finkel, D. Korytov, E. Kovacs, S. Rizzi, and J. Insley. "The Last Journey. I. An Extreme-Scale Simulation on the Mira Supercomputer," *The Astrophysical Journal Supplement Series* (January 2021), IOP Publishing.

Nucleon Axial Charge with All-Staggered Lattice QCD

PI Andreas Kronfeld, Fermi National Accelerator Laboratory

AWARD ALCC


HOURS Theta: 200,000 Node-Hours

With help from DOE supercomputing resources, researchers are performing lattice quantum chromodynamics (QCD) calculations to investigate physics beyond the Standard Model and inform several major experimental programs in high energy and nuclear physics, including CERN's Large Hadron Collider and Fermilab's upcoming Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE).

CHALLENGE As neutrino-oscillation experiments move from first observations to precise measurements, the theoretical understanding of the neutrino-nucleus interaction will play an ever more important role. The nucleon axial form factor, which parametrizes the weak responses of a proton or neutron, is difficult to measure experimentally, and it is a dominant uncertainty in neutrino-nuclei cross-section calculations for incoming neutrino energies at around 1 GeV. Therefore, a theoretical calculation with lattice QCD provides a non-ambiguous determination of the form factor that could help reducing the uncertainty.

APPROACH With this ALCC project, the team is taking the next step in demonstrating that staggered valence quarks are a viable strategy in lattice QCD for nucleon physics. They are using staggered valence quarks on ensembles of lattice gauge field configurations from the MILC Collaboration to compute the nucleon axial charge, a hadronic matrix element entering the neutron decay rate and, simultaneously, the normalization of the nucleon axial form factor. The team's software stack is a branch of the MILC codes extended to baryon two- and three-point lattice QCD correlation functions.

RESULTS The researchers are using the highly improved staggered quark (HISQ) action for both valence and sea



Left: Current interplay between data, hadronic physics (circle enclosing a red, green, and blue quark), and nuclear physics (collection of nucleons at the bottom), feeding into the GENIE Monte Carlo, and ultimately uncertainties on oscillation parameters, such as θ_{23} . Right: Sketch of the disruption possible with ab initio lattice QCD calculations. Image: Andreas Kronfeld, Fermilab

quarks. They are employing the library of 2+1+1-flavor ensembles of gluon fields from the MILC Collaboration. In this way, they are able to attain high statistics and work directly at the light quark mass corresponding to the average of the up- and down-quark masses in nature. Building on previous work, the team is using ALCF and NERSC supercomputers to generate data at additional lattice spacings, $a \approx 0.09$ fm and 0.06 fm, thereby enabling a robust, nontrivial extrapolation to the continuum limit.

IMPACT The team's calculations of the axial charge are vital to DOE's High Energy Physics and Nuclear Physics programs. The calculations are a first step toward creating a suite of lattice QCD computations that target nucleon matrix elements and are essential ingredients for building the nuclear models needed by neutrino-oscillation long baseline experiments. In addition, the lattice QCD calculations can provide insights into matrix elements that are inaccessible to experiment.

Physics |  Simulation

Amplification of Cosmic Magnetic Fields Captured in Laser Experiments

PI Petros Tzeferacos, University of Rochester


AWARD ALCC

HOURS Mira: 1,375,000 Node-Hours

The origin of cosmic magnetic fields that thread our universe is not fully understood. Astrophysicists believe that a key mechanism behind their origin is turbulent dynamo, a process by which stochastic motions amplify and sustain cosmic magnetic fields at the values we observe today. Although this physical process occurs readily in space, it is extremely hard to recreate and study in the laboratory. Researchers of the turbulent dynamo collaboration, an international team co-led by the University of Rochester and the University of Oxford performed laser experiments that captured—for the first time in the laboratory—the growth rate of magnetic fields by the turbulent dynamo mechanism, in conditions relevant to most astrophysical systems.

CHALLENGE These laser-driven plasma experiments were able to reproduce experimentally turbulent dynamo and, for the first time in the laboratory, access the viscosity-dominated regime that is relevant to most plasmas in the universe. The scientists were also able, for the first time, to record time-resolved measurements of the properties of the mechanism, including the growth rate of the magnetic field, previously only available from simulations.

APPROACH To design their experiments (conducted at the University of Rochester's Omega Laser Facility), the researchers leveraged Mira to construct a novel experimental platform using numerical simulations performed with FLASH, a publicly available simulation code that can accurately model laser-driven laboratory plasma experiments. The platform consists of a pair of plastic foils that are driven with twenty OMEGA laser beams. The laser ablation launches a pair of magnetized plasma flows that propagate through offset grids, and collide and shear to create a hot, turbulent plasma. The turbulent plasma achieves a regime where



FLASH simulation of the experiment showing the magnetic fields achieved.
Image: Petros Tzeferacos, University of Rochester

turbulent dynamo can amplify the advected seed magnetic fields to magnetic energies comparable to the kinetic energy of the stochastic motions.

RESULTS As detailed in a paper published in *Proceedings of the National Academy of Sciences*, the researchers found rapid magnetic field amplification that exceeds theoretical expectations, a result that could help explain the origin of present-day large-scale fields observed in galaxy clusters.

IMPACT The ability to recreate and study turbulent dynamo in the laboratory is a long-sought goal for physicists and astrophysicists. The properties of this fundamental mechanism were previously only available from simulations and analytical calculations. These experiments now provide strict constraints to existing models and help explain astronomical measurements. The efficient amplification of magnetic fields at large scales seen in the experiments could explain the origin of large-scale fields observed in galaxy clusters, which are not predicted by current idealized simulations. These experiments answer key astrophysics questions and establish laboratory experiments as a component in the study of turbulent dynamo.

PUBLICATIONS

Bott, A. F. A., P. Tzeferacos, L. Chen, C. A. J. Palmer, A. Rigby, A. R. Bell, R. Bingham, A. Birkel, C. Graziani, D. H. Froula, J. Katz, M. Koenig, M. W. Kunz, C. Li, J. Meinecke, F. Miniati, R. Petrasso, H.-S. Park, B. A. Remington, B. Reville, J. S. Ross, D. Ryu, D. Ryutov, F. H. Séguin, T. G. White, A. A. Schekochihin, D. Q. Lamb, and G. Gregori. "Time-Resolved Turbulent Dynamo in a Laser Plasma," *Proceedings of the National Academy of Sciences* (March 2021), National Academy of Sciences.

Towards a Definitive Model of Core-Collapse Supernova Explosions

PI Adam Burrows, Princeton University

AWARD INCITE


HOURS Theta: 2,000,000 Node-Hours

Core-collapse supernovae, which produce the highest densities of matter and energy in the universe, are responsible for creating most of the elements in nature. To shed light on this mysterious cosmological phenomenon, a research team led by Princeton University is using ALCF supercomputers to complete and analyze the largest-ever collection of 3D supernovae simulations, spanning a broad range of progenitor masses and structures.

CHALLENGE Since the 1960s, there has been an agonizingly slow march towards demonstrating a robust mechanism of supernovae explosion. 2D simulations of supernovae have supported the theory that capturing a small fraction of the neutrinos emitted during collapse powers the explosions, but detailed 3D calculations proving this paradigm were lacking. With the power of leadership-class supercomputers and continued advances in software, researchers now have the capabilities to tackle this longstanding challenge in nuclear astrophysics.

APPROACH To carry out their studies, the team is using FORNAX, their highly scalable, 3D radiation-hydrodynamics code. The team is running large-scale simulations on Theta aimed at determining if the neutrino mechanism is a robust explanation for supernova explosions and the stellar progenitor dependence of their observables.

RESULTS In a paper published in *Nature*, the team reviewed the status of supernova theory and drew insights from eight models generated on Theta. One of the most important conclusions of their studies is that the most massive progenitor models need to be continued for longer physical times, perhaps to many seconds, to asymptote to a final state, in particular vis à vis explosion energy. Those models that explode more lethargically and a bit later after



This visualization depicts the inner turbulent convective region 146 ms after core bounce and before the explosion, showing accreted matter tracers swirling randomly about the newly birthed proto-neutron star core. *Image: Joseph A. Insley, Argonne National Laboratory; Adam Burrows, Princeton University and David Vartanyan, University of Berkeley, California.*

bounce tend not to include much neutron-rich ejecta, while those that explode more quickly, such as the lowest-mass progenitors, can eject some more neutron-rich matter. However, in all their 3D models, the inner ejecta has a net proton-richness. If true, this systematic result has important consequences for the nucleosynthesis yields as a function of progenitor.

IMPACT The team's efforts to advance the fundamental theoretical understanding of supernova explosions will benefit ongoing research efforts to determine the origin of the elements in the universe, measure gravitational waves, and interpret laboratory nuclear reaction rate measurements in light of stellar nucleosynthesis.

PUBLICATIONS

A. Burrows and D. Vartanyan, "Core-Collapse Supernova Explosion Theory", *Nature*, (January 2021).

H. Nagakura, A. Burrows, D. Vartanyan, and D. Radice, "Core-collapse supernova neutrino emission and detection informed by state-of-the-art three-dimensional numerical models," *Monthly Notices of the Royal Astronomical Society* (October 2020), Oxford University Press.

ALCF Projects

2021 INCITE

CHEMISTRY

Design of Peptides and Proteins on Classical and Quantum Computing Hardware

PI Vikram Mulligan, Flatiron Institute
HOURS ALCF: 800,000 Node-Hours

EARTH SCIENCE

Energy Exascale Earth System Model

PI Mark Taylor, Sandia National Laboratories
HOURS ALCF: 1,800,000 Node-Hours
OLCF: 1,000,000 Node-Hours

Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling

PI Christine Goulet, University of Southern California
HOURS ALCF: 480,000 Node-Hours
OLCF: 319,000 Node-Hours

ENERGY TECHNOLOGIES

DNS and LES of Internal Combustion Engines to Understand Origins of CCV

PI Sibendu Som, Argonne National Laboratory
HOURS ALCF: 1,000,000 Node-Hours

ENGINEERING

Adaptive DDES of a Vertical Tail/Rudder Assembly with Active Flow Control

PI Kenneth Jansen, University of Colorado Boulder
HOURS ALCF: 1,000,000 Node-Hours

High-Speed Turbulence with Shocks over Non-Adiabatic and Flexible Walls

PI Johan Larsson, University of Maryland
HOURS ALCF: 1,800,000 Node-Hours

MATERIALS SCIENCE

First-Principles Simulation of Electronic Stopping Excitation and Beyond

PI Yosuke Kanai, University of North Carolina at Chapel Hill
HOURS ALCF: 1,900,000 Node-Hours

Large-Scale Simulations of Light-Activated Matter

PI Giulia Galli, University of Chicago and Argonne National Laboratory
HOURS ALCF: 1,200,000 Node-Hours

Predicting Ion Transport Kinetics at Complex Interfaces for Energy Storage

PI Brandon Wood, Lawrence Livermore National Laboratory
HOURS ALCF: 640,000 Node-Hours

Towards a Definitive Model of Core-Collapse Supernova Explosions

PI Adam Burrows, Princeton University
HOURS ALCF: 2,000,000 Node-Hours

Towards Predictive Simulations of Functional and Quantum Materials

PI Paul Kent, Oak Ridge National Laboratory
HOURS ALCF: 1,800,000 Node-Hours
OLCF: 500,000 Node-Hours

Ultrafast Control of Functional Materials

PI Priya Vashishta, University of Southern California
HOURS ALCF: 1,600,000 Node-Hours

PHYSICS

Ab-initio Nuclear Structure and Nuclear Reactions

PI Gaute Hagen, Oak Ridge National Laboratory
HOURS ALCF: 500,000 Node-Hours
OLCF: 693,000 Node-Hours

High-Fidelity Gyrokinetic Simulation of Tokamak and ITER Edge Physics

PI Choongseock Chang, Princeton Plasma Physics Laboratory
HOURS ALCF: 1,300,000 Node-Hours
OLCF: 900,000 Node-Hours

Petascale Simulations of Kinetic Effects in IFE Plasmas

PI Frank Tsung, University of California, Los Angeles
HOURS ALCF: 1,800,000 Node-Hours
OLCF: 900,000 Node-Hours

Towards a Definitive Model of Core-Collapse Supernova Explosions

PI Adam Burrows, Princeton University
HOURS Theta: 2,000,000 Node-Hours

ALCC 2020–2021

COMPUTER SCIENCE

Enabling Resilient and Portable Workflows from DOE's Experimental Facilities

PI Katie Antypas, Lawrence Berkeley National Laboratory
HOURS ALCF: 100,000 Node-Hours
OLCF: 20,000 Node-Hours

CHEMISTRY

Benchmarking Many-Body Perturbation Theory

PI Olle Heinonen, Argonne National Laboratory
HOURS ALCF: 100,000 Node-Hours

Interpretable Machine Learning Force Fields for Accurate Chemical Reactive Dynamic

PI Olexandr Isayev, Carnegie Mellon University
HOURS ALCF: 359,000 Node-Hours

Stochastic A Priori Dynamics for Complex Reactive Chemical Environments

PI Ahren Jasper, Argonne National Laboratory
HOURS ALCF: 100,000 Node-Hours

Understanding the Role of Hierarchical Correlations in Solution-Phase Chemical Separations

PI Lynda Soderholm, Argonne National Laboratory
HOURS ALCF: 100,000 Node-Hours

EARTH SCIENCE

Variable-Resolution Earth System Modeling of the Cryosphere with E3SM

PI Darin Comeau, Los Alamos National Laboratory
HOURS ALCF: 400,000 Node-Hours
NERSC: 500,000 Node-Hours

ENERGY TECHNOLOGIES

Automatic Building Energy Modeling

PI Joshua New, Oak Ridge National Laboratory
HOURS ALCF: 300,000 Node-Hours

DNS Simulations of Coolant Flow in the High-Flux Isotope Reactor

PI Emilian Popov, Oak Ridge National Laboratory
HOURS ALCF: 220,000 Node-Hours

Multiphase Flow Simulations of Reactor Flows

PI Igor Bolotnov, North Carolina State University
HOURS ALCF: 192,000 Node-Hours
NERSC: 250,000 Node-Hours

Toward the Future: High-Fidelity Simulation for Next-Generation Nuclear Reactors

PI Yiqi Yu, Argonne National Laboratory
HOURS ALCF: 208,000 Node-Hours
NERSC: 300,000 Node-Hours

MATERIALS SCIENCE

Data-Driven Molecular Engineering of Advanced Functional Materials

PI Jacqueline Cole, University of Cambridge
HOURS ALCF: 100,000 Node-Hours

High-Temperature Material Properties from First Principles

PI Mark Messner, Argonne National Laboratory
HOURS ALCF: 200,000 Node-Hours

Many-Body Perturbation Theory Meets Machine Learning to Discover Materials for Organic Photovoltaics

PI Noa Marom, Carnegie Mellon University
HOURS ALCF: 100,000 Node-Hours

Plasma Surface Interaction Modeling

PI Brian Wirth, University of Tennessee
HOURS ALCF: 318,000 Node-Hours
OLCF: 155,000 Node-Hours
NERSC: 30,000 Node-Hours

Predictive Modeling of Nanoporous Materials and Multiphase Systems

PI Joern Siepmann, University of Minnesota
HOURS ALCF: 220,000 Node-Hours

Supercomputing for Automotive High-Temperature Alloy Design

PI Dongwon Shin, Oak Ridge National Laboratory
HOURS ALCF: 100,000 Node-Hours

PHYSICS

Reconstructing Neutrino Data with the MicroBooNE Liquid Argon Detector

PI Andrzej Szelc, University of Manchester
HOURS ALCF: 200,000 Node-Hours

Optimization Studies of the LBNF - PIP-II Complex for Megawatt Beams on Target

PI Igor Rakhno, Fermi National Accelerator Laboratory
HOURS ALCF: 450,000 Node-Hours

Chiral Nuclear Interactions from Nuclei to Nucleonic Matter

PI Maria Piarulli, Washington University in St. Louis
HOURS ALCF: 200,000 Node-Hours

Nucleon Axial Charge with All-Staggered Lattice QCD

PI Andreas Kronfeld, Fermi National Accelerator Laboratory
HOURS ALCF: 200,000 Node-Hours
NERSC: 870,000 Node-Hours

Distributed Large Wavefield Propagation and 3D Reconstruction Beyond the Depth of Focus Limit

PI Ming Du, Argonne National Laboratory
HOURS ALCF: 250,000 Node-Hours

Field-Reversed Configuration Stability and Transport

PI Sean Dettrick, TAE Technologies Inc.
HOURS ALCF: 64,000 Node-Hours

Improving Direct-Drive Laser Fusion Predictive Capability with a Simulation Database

PI Duc Cao, Laboratory for Laser Energetics
HOURS ALCF: 318,000 Node-Hours

Improving Direct-Drive Laser Fusion Predictive Capability with a Simulation Database

PI Duc Cao, Laboratory for Laser Energetics
HOURS ALCF: 318,000 Node-Hours

High-Luminosity LHC Detector Upgrade Studies by the ATLAS and CMS Collaborations

PI Douglas Benjamin, Argonne National Laboratory
HOURS ALCF: 950,000 Node-Hours

ALCC 2021–2022

BIOLOGICAL SCIENCES

Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

PI Jonathan Ozik, Argonne National Laboratory
HOURS ALCF: 160,000 Node-Hours

CHEMISTRY

Microscopic Insight Into Transport Properties of Li-Battery Electrolytes

PI Wei Jiang, Argonne National Laboratory
HOURS ALCF: 1,032,000 Node-Hours

Multimodal Imaging with Intense X-ray Pulses

PI Phay Ho, Argonne National Laboratory
HOURS ALCF: 316,000 Node-Hours

EARTH SCIENCE

Advancing Multi-Year to Decadal Climate Prediction with High-Resolution E3SM and CESM

PI Ben Kirtman, University of Miami
HOURS ALCF: 1,000,000 Node-Hours

Improving the Representation of Mesoscale Convective Systems in Weather and Climate

PI Andreas Prein, National Center for Atmospheric Research
HOURS ALCF: 372,000 Node-Hours

Multi-Decadal, Climate-Scale Convection-Resolving Simulations for North America

PI V. Rao Kotamarthi, Argonne National Laboratory
HOURS ALCF: 300,000 Node-Hours

Multi-Scale Multi-Physics Ensemble Simulations for Aerosol-Cloud Interactions

PI Po-Lun Ma, Pacific Northwest National Laboratory
HOURS ALCF: 400,000 Node-Hours

ENERGY TECHNOLOGIES

Cavitation Dynamics in the Spallation Neutron Source Target

PI Eric Johnsen, University of Michigan
 HOURS ALCF: 108,000 Node-Hours
 OLCF: 4,000 Node-Hours

High-Fidelity CFD Simulations Supporting the Needs of Industry and the DOE

PI Dillon Shaver, Argonne National Laboratory
 HOURS ALCF: 550,000 Node-Hours
 OLCF: 80,000 Node-Hours

High-Fidelity Physics Simulations for DOE and Industry Fast Spectrum Nuclear Reactor Systems

PI Emily Shemon, Argonne National Laboratory
 HOURS ALCF: 880,000 Node-Hours

High-Fidelity Simulation of Flow and Heat Transfer Behavior to Support Conversion of Research Reactors with Involute Shaped Fuel Elements to Low Enriched Uranium

PI Yiqi Yu, Argonne National Laboratory
 HOURS ALCF: 500,000 Node-Hours

ENGINEERING

Multiscale Bubble Breakup and Gas Transfer in Turbulent Oceanic Environments

PI Parviz Moin, Stanford University
 HOURS ALCF: 650,000 Node-Hours
 NERSC: 500,000 Node-Hours

MATERIALS SCIENCE

Inverse Design of Multicomponent Oxide Catalysts with Generative Models and DFT

PI Rafael Gomez-Bombarelli, Massachusetts Institute of Technology
 HOURS ALCF: 400,000 Node-Hours
 OLCF: 50,000 Node-Hours
 NERSC: 200,000 Node-Hours

Modeling of Polymeric Materials for Energy Storage Across Scales

PI Juan de Pablo, University of Chicago
 HOURS ALCF: 100,000 Node-Hours

Response Functions of LaNiO₂: Insights into High-Temperature Superconductivity

PI Gabriel Kotliar, Rutgers University
 HOURS ALCF: 115,000 Node-Hours

PHYSICS

AI/Deep Learning Prediction & Real-Time Control for Fusion Energy Systems

PI William Tang, Princeton University
 HOURS ALCF: 45,000 Node-Hours

LBNF - PIP-II Optimization Studies for Megawatt 120-GeV Beams on Target

PI Igor Rakhno, Fermi National Accelerator Laboratory
 HOURS ALCF: 450,000 Node-Hours

Quantum Monte Carlo Calculations of Nuclei Up to ¹⁶O and Neutron matter

PI Alessandro Lovato, Argonne National Laboratory
 HOURS ALCF: 632,000 Node-Hours

ALCF DATA SCIENCE PROGRAM

Advanced Materials Characterization with AI-Informed Computation

PI Marco Govoni, Argonne National Laboratory

Autonomous Molecular Design for Redox Flow Batteries

PI Logan Ward, Argonne National Laboratory

Deep Learning at Scale for Multimessenger Astrophysics Through the NCSA-Argonne Collaboration

PI Eliu Huerta, University of Illinois at Urbana-Champaign

Developing High-Performance-Computing Applications for Liquid Argon Neutrino Detectors

PI Andrzej Szcel, The University of Manchester

Dynamic Compressed Sensing for Real-Time Tomographic Reconstruction

PI Robert Hovden, University of Michigan

Learning Optimal Image Representations for Current and Future Sky Surveys

PI George Stein, Lawrence Berkeley National Laboratory

Machine Learning for Data Reconstruction to Accelerate Physics Discoveries in Accelerator-Based Neutrino Oscillation Experiments

PI Marco Del Tutto, Fermi National Accelerator Laboratory

Machine Learning Magnetic Properties of Van Der Waals Heterostructures

PI Trevor Rhone, Rensselaer Polytechnic Institute

X-ray Microscopy of Extended 3D Objects: Scaling Towards the Future

PI Chris Jacobsen, Argonne National Laboratory and Northwestern University

AURORA EARLY SCIENCE PROGRAM

Accelerated Deep Learning Discovery in Fusion Energy Science

PI William Tang, Princeton Plasma Physics Laboratory

Dark Sky Mining

PI Salman Habib, Argonne National Laboratory

Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics

PI Kenneth Jansen, University of Colorado Boulder

Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

PI Nicola Ferrier, Argonne National Laboratory

Exascale Computational Catalysis

PI David Bross, Argonne National Laboratory

Extending Moore's Law Computing with Quantum Monte Carlo

PI Anouar Benali, Argonne National Laboratory

Extreme-Scale Cosmological Hydrodynamics

PI Katrin Heitmann, Argonne National Laboratory

Extreme-Scale In-Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations

PI Amanda Randles, Duke University and Oak Ridge National Laboratory

Extreme-Scale Unstructured Adaptive CFD

PI Kenneth Jansen, University of Colorado Boulder

High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

PI C.S. Chang, Princeton Plasma Physics Laboratory

Machine Learning for Lattice Quantum Chromodynamics

PI William Detmold, Massachusetts Institute of Technology

Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials

PI Noa Marom, Carnegie Mellon University

NWChemEx: Tackling Chemical, Materials, and Biochemical Challenges in the Exascale Era

PI Thom Dunning, Pacific Northwest National Laboratory

Simulating and Learning in the ATLAS Detector at the Exascale

PI Walter Hopkins, Argonne National Laboratory

Virtual Drug Response Prediction

PI Rick Stevens, Argonne National Laboratory

2020 DIRECTOR'S DISCRETIONARY

The following list provides a sampling of the many Director's Discretionary projects at the ALCF.

BIOLOGICAL SCIENCES

AI-Driven Drug Discovery for SARS-CoV-2 Proteome

PI Arvind Ramanathan, Argonne National Laboratory

Computational Physical Genomics: Exploring Potential Novel Cancer Therapies

PI Allen Taflove, Northwestern University

COVID-19 Spread and Effectiveness of Interventions

PI Jonathan Ozik, Argonne National Laboratory

Finding Druggable Sites in SARS-CoV-2 Proteins Using Molecular Dynamics and Machine Learning

PI Albert Lau, John Hopkins School of Medicine

Modeling Coronavirus

PI Zhangli Peng, University of Illinois at Chicago

Real-Time Analysis of SARS-CoV-2 Proteins

PI Darren Sherrell, Argonne National Laboratory

Simulation of Viral Infection Propagation Through Air-Travel

PI Ashok Srinivasan, University of West Florida

CHEMISTRY

Enantioselectivity in Heterogeneous Catalysts via the Addition of Chiral Modifiers

PI Wilfred T. Tysoe, University of Wisconsin-Milwaukee

Simulation of Industrial Flares

PI Marc Cremer, Reaction Engineering International

COMPUTER SCIENCE

DLIO: A Data-Centric Benchmark for Scientific Deep Learning Applications

PI Huihuo Zheng and Venkatram Vishwanath, Argonne National Laboratory

MPICH - A High Performance and Widely Portable MPI Implementation

PI Ken Raffenetti, Argonne National Laboratory

Rendezvous Algorithms for Large-Scale Modeling and Simulation

PI Steven Plimpton, Sandia National Laboratories

EARTH SCIENCE

Generation of a Next-Level Dataset for Regional Scale Climate Modeling: Convective Resolving Spatial Scales

PI Jiali Wang, Argonne National Laboratory

Using ARM Observations to Evaluate Process-Interactions in MCS Simulations Across Scales

PI Andreas Franz Prein, National Center for Atmospheric Research

ENERGY TECHNOLOGIES

Automatic Building Energy Modeling and Analysis

PI Joshua Ryan New, Oak Ridge National Laboratory Energy Technologies

Cask Mis-Loads Evaluation Techniques

PI Angela Di Fulvio, University of Illinois at Urbana-Champaign

High-Fidelity CFD Simulations of Multi-Mode Combustion

PI Pinaki Pal, Argonne National Laboratory

LES Simulations of Severe Accident Conditions in Nuclear Containment

PI Christopher Boyd, Westinghouse Electric Company

ENGINEERING

Nek5000/NekRS for COVID Particle Transport and NRC

PI Aleksandr V. Obabko, Argonne National Laboratory

Towards Reactive DNS in Complex Internal Combustion Engine Geometries

PI Christos Frouzakis, ETH Zurich

MATERIALS SCIENCE

Designing Catalysts under Realistic Operating Conditions from the Atomic Scale

PI Hendrik Heinz, University of Colorado Boulder

Materials Informatics Study of Two-Dimensional Magnetic Materials and Their Heterostructures

PI Trevor Rhone, Rensselaer Polytechnic Institute

Structure and Properties of Grain Boundaries in Materials for Energy Applications

PI Wissam A Saidi, University of Pittsburgh

PHYSICS

Experimentally Driven Automated Machine-Learned Interatomic Potential for a Refractory Oxide

PI Ganesh Sivaraman, Argonne National Laboratory

LatticeQCD: Exascale Lattice Gauge Theory Opportunities/Requirements for Nuclear & High Energy Physics

PI Andreas Kronfeld, Fermilab

About the Argonne Leadership Computing Facility

Argonne's Leadership Computing Facility Division operates the Argonne Leadership Computing Facility (ALCF) as part of the U.S. Department of Energy's effort to provide leadership-class computing resources to the scientific community. The ALCF is supported by the DOE Office of Science, Advanced Scientific Computing Research (ASCR) program.

About Argonne National Laboratory

Argonne is a U.S. Department of Energy Laboratory managed by UChicago Argonne, LLC, under contract DE-AC02-06CH11357. The Laboratory's main facility is outside of Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, visit www.anl.gov.

Availability of this Report (ANL/ALCF-21/3)

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available for free via DOE's SciTech Connect (<http://www.osti.gov/scitech/>).

Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd.
Alexandria, VA 22312

phone | 800.553.NTIS (6847) or 703.605.6000
fax | 703.605.6900
orders@ntis.gov
www.ntis.gov

Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062

phone | 865.576.8401
fax | 865.576.5728
reports@adonis.osti.gov
www.osti.gov

ALCF Leadership: Michael E. Papka (Division Director), Jini Ramprakash (Deputy Division Director), Katherine Riley (Director of Science), Kalyan Kumaran (Director of Technology), Mark Fahey (Director of Operations), and Susan Coghlan (ALCF-X Project Director)

Editorial Team: Beth Cerny, Jim Collins, Nils Heinonen, Logan Ludwig, and Laura Wolf

Design and production: Sandbox Studio, Chicago

Disclaimer: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.



CONTACT

media@alcf.anl.gov

alcf.anl.gov



Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

